Selected Subsets of Information for

Meeting With Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm
GSFC building 33, room H114

J. Patrick (Pat) Gary
Network Projects Leader
Networks and Information Technology Security Group (Code 606.1)
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Agenda

- Welcome +
  Meeting Purpose — Pat Gary
- Intro of HECN+NREN Network Staff +
  Intro to Network-Performance-Affecting Factors — Pat Gary
- NREN-initiated User-Oriented Performance-Affecting Efforts — Ken Freeman
- Technical Examples of Performance-Affecting Efforts — Dave Hartzell
- HECN-initiated User-Oriented Performance-Affecting Efforts — Pat Gary et al
- Q&A +
  Next Steps — All

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Meeting Purpose

- Improved awareness for getting network-related assistance
- A brief overview of a range of ARC/NAS/NREN staff efforts to help Columbia users and their SA's improve the throughput performance of their file transfers with Columbia
- Technical examples of efforts which have significantly improved the throughput performance of file transfers with Columbia
- Opportunities to schedule/plan more specific meetings with NCCS+Columbia networking staff/experts

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Agenda

- Welcome +
  Meeting Purpose — Pat Gary
- Intro of HECN+NREN Network Staff +
  Intro to Network-Performance-Affecting Factors — Pat Gary
- NREN-initiated User-Oriented Performance-Affecting Efforts — Ken Freeman
- Technical Examples of Performance-Affecting Efforts — Dave Hartzell
- HECN-initiated User-Oriented Performance-Affecting Efforts — Pat Gary et al
- Q&A +
  Next Steps — All

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Intro to Network-Performance-Affecting Factors
Note: NOT even close to a complete list!

- User Computer Based
  - File transfer application + disk I/O speeds
  - Host computer speed
  - TCP tuning
  - Network Interface Card (inc. being jumbo frame capable)
- LAN Based (inc. being jumbo frame capable end-to-end)
  - Connection infrastructure
  - Perimeter firewall
- WAN Based (inc. being jumbo frame capable end-to-end)
  - Round Trip Time (RTT)
  - Packet loss rate

On TCP Dynamics

- **TCP Window Buffer Size (bytes) = Bandwidth (bps) * Delay (RTT seconds)**

- **The Mathis Equation for Reno Paraphased**

$$Speed \leq \frac{MTU}{RTT * \sqrt{loss}}$$

  - M. Mathis et al, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", Computer Communications Review, volume 27, number 3, July 1997

Internet2 Land Speed Record

(Rules and current records: http://lsr.internet2.edu/)

Last IPv4 Single Stream Record (http://data-reservoir.adm.s.u-tokyo.ac.jp/lsr-20060219/)

- **8.80 Gbps** (TCP payload), 9014B frame, 264,147 terabit-meters / second
- 32,372 km connection from Tokyo through Seattle, Amsterdam, Chicago to Tokyo Latency 500ms RTT

NETWORK USED IN THE EXPERIMENT: nearly three-quarters the circumference of the Earth

On February 3, 2006, GSFC's Bill Fink simultaneously conducted two 6-hour-duration nuttcp-enabled UDP-based 4.6-Gbps flow tests across the NLR-based 10-Gbps NREN lambda . For this test the transmit and receive paths of the NREN lambda were looped at Sunnyvale, causing the test data to flow in both directions. The tests filled the 10-Gbps NREN lambda to 92% of capacity and recorded no packet losses.

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

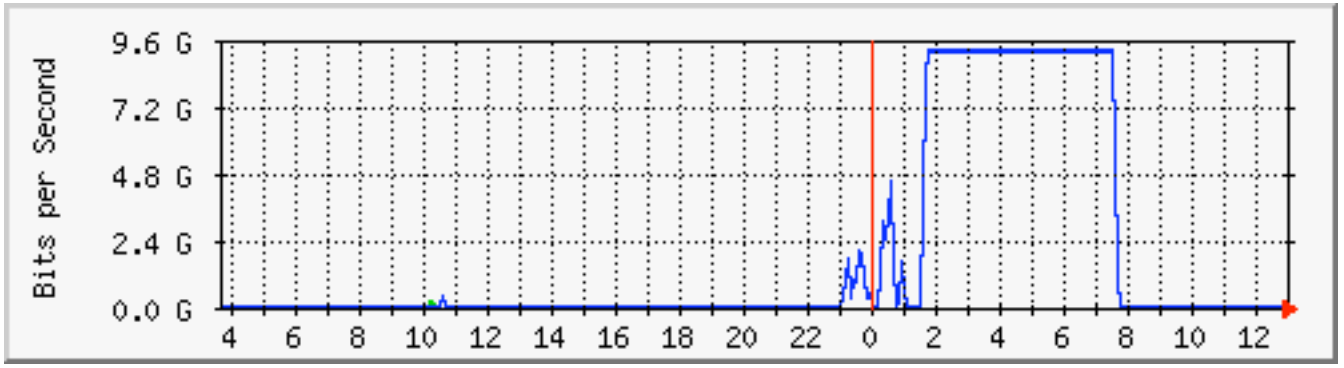

NLR/NREN TX 10 GigE
3 February 2006

Max In:	350.4 Mb/s (3.5%)	Average In:	1879.9 kb/s (0.0%)	Current In:	0.0 b/s (0.0%)
Max Out:	9305.3 Mb/s (93.1%)	Average Out:	1779.1 Mb/s (17.8%)	Current Out:	0.0 b/s (0.0%)

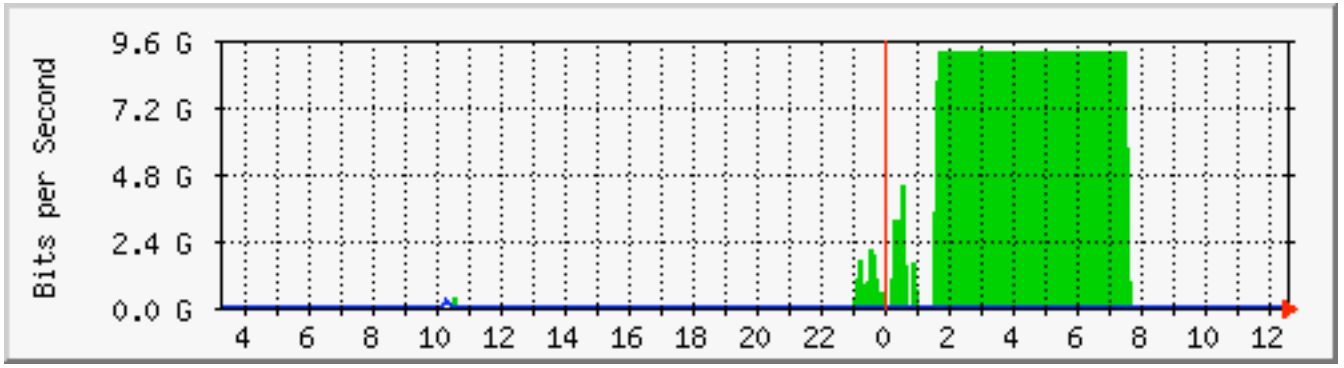

NLR/NREN RX 10 GigE
3 February 2006

Max In:	9316.8 Mb/s (93.2%)	Average In:	1779.1 Mb/s (17.8%)	Current In:	0.0 b/s (0.0%)
Max Out:	298.9 Mb/s (3.0%)	Average Out:	1879.9 kb/s (0.0%)	Current Out:	0.0 b/s (0.0%)

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Agenda

- Welcome +
  Meeting Purpose — Pat Gary
- Intro of HECN+NREN Network Staff +
  Intro to Network-Performance-Affecting Factors — Pat Gary
- NREN-initiated User-Oriented Performance-Affecting Efforts — Ken Freeman
- Technical Examples of Performance-Affecting Efforts — Dave Hartzell
- HECN-initiated User-Oriented Performance-Affecting Efforts — Pat Gary et al
- Q&A +
  Next Steps — All

GSFC SEN and HECN

Service Description Overview

- **Scientific and Engineering Network (SEN),** typically enabling 1-10 gigabit per second (Gbps) user connections: a non-mission-dedicated high-end computer network at GSFC serving GSFC projects/users who have computer network performance requirements greater than those baselined for GSFC's general-use campus-wide Center Network Environment (CNE)
  - http://cisto.gsfc.nasa.gov/SEN.html
- **High End Computer Network (HECN)/Lambda Network (L-Net),** typically enabling 10+ Gbps R&D connections: network R&D and testbed evaluations with advanced network technology to contribute to the next generation high-end computer networks at GSFC
  - http://cisto.gsfc.nasa.gov/IRAD_Lambda.html
- Both managed by Code 606.1's HECN Team

GSFC SEN and HECN

SEN+HECN User Services

- Local Connectivity/Access
- Remote Access
- Production IP Services
- Testbed for Advanced IP Services
- Help Desk

Notional Key Characteristics of GSFC's Scientific and Engineering Network (SEN) and High End Computer Network (HECN/L-Net

[Note: Some data flows/paths are restricted by GSFC security policy and/or management agreement]

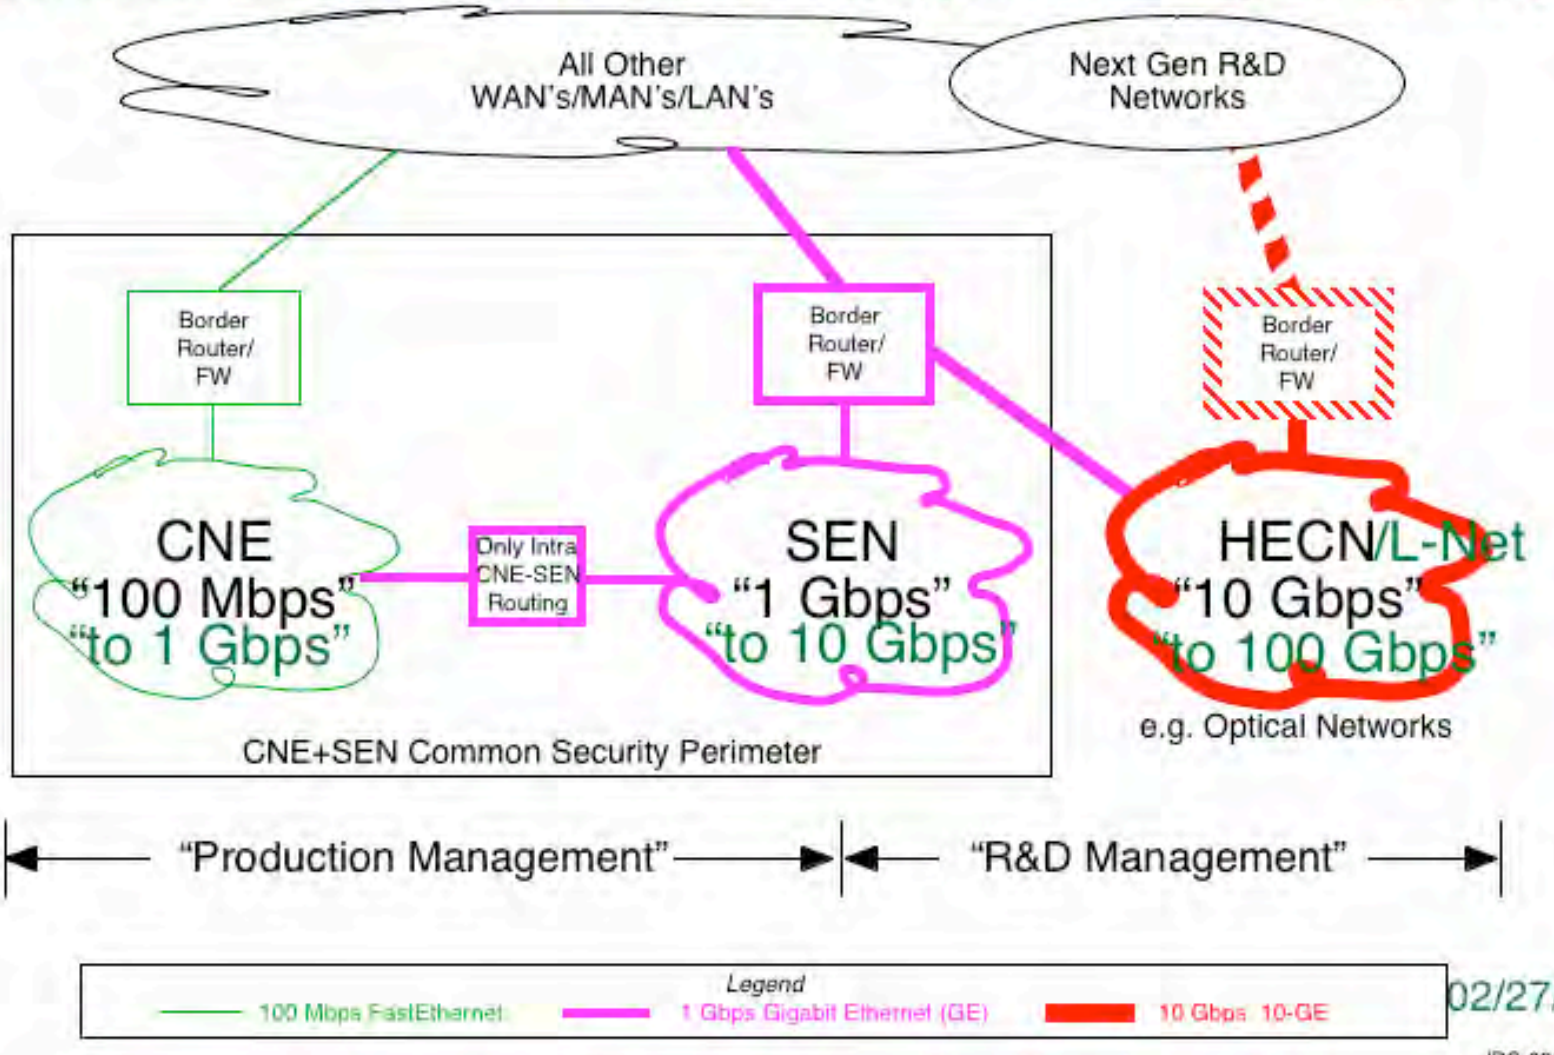

02/27/06

JPG 02/25/04

GSFC Scientific and Engineering Network (SEN) Major Links

Circa 1 ~~March~~ October 2006

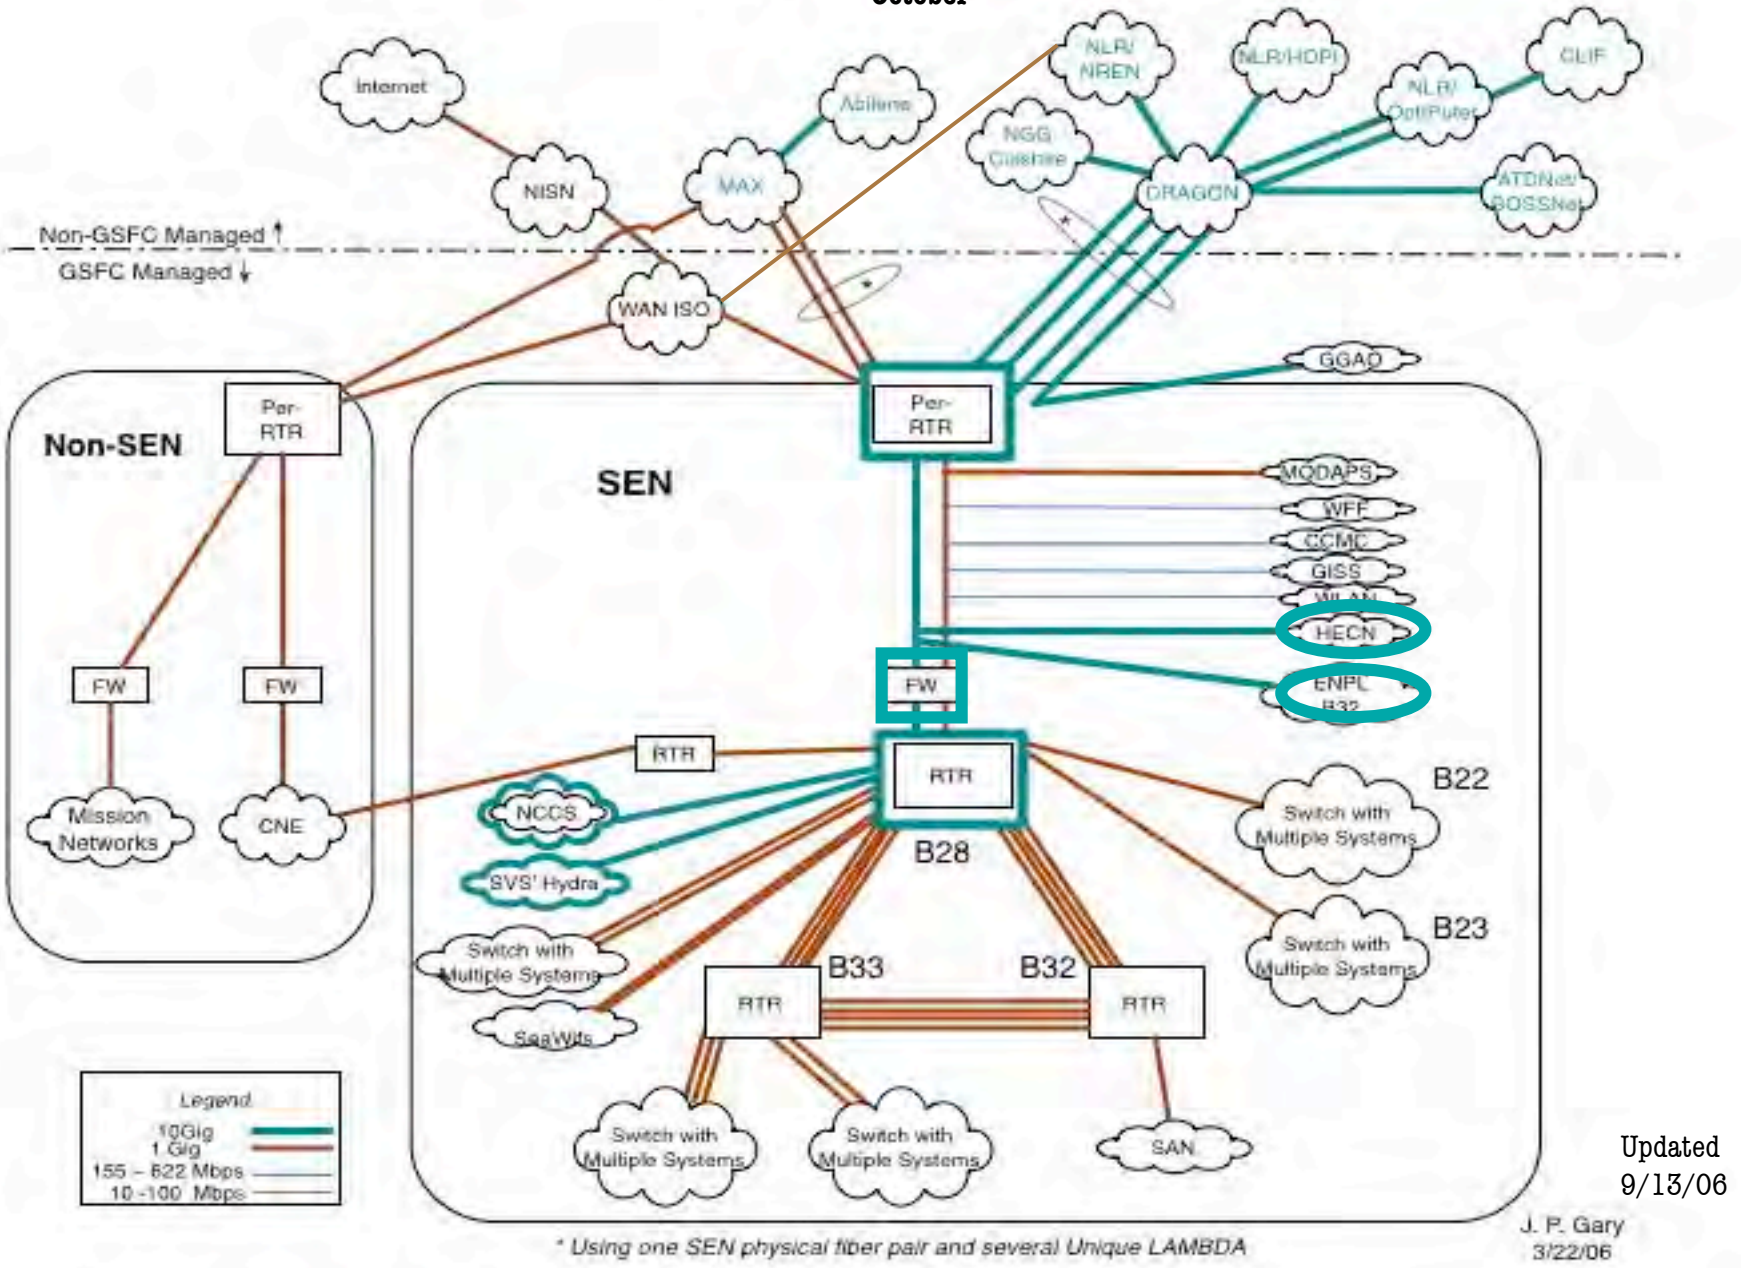

* Using one SEN physical fiber pair and several Unique LAMBDA

J. P. Gary
3/22/06

Updated
9/13/06

GSFC SEN and HECN

http://cisto.gsfc.nasa.gov/SENuserdocs/SENuser.html

- SEN Overview
- SEN Host Registration Information
- SEN Firewall Information
- Throughput Performance Tuning Information
- Recent SEN Alert/Notices
- Requesting SEN Troubleshooting
- Related Links
- Contact Us

GSFC SEN and HECN

Throughput Performance Tuning Information

- System Specific Notes for System Administrators (and Privileged Users): http://www.psc.edu/networking/projects/tcptune/
- TCP Tuning Guide: http://www-didc.lbl.gov/TCP-tuning
- Phil Dykstra's tutorial: http://www.wcisd.hpc.mil/%7Ephil/sc2004
- M. Mathis, et al, "NPAD/pathdiag unleashed", ESCC/Internet2 Joint Techs Workshop, Madison, July, 2006: http://events.internet2.edu/2006/jt-madison/sessionDetails.cfm?session=2753&event=253
- SEN Overview, especially pp. 40-41: http://cisto.gsfc.nasa.gov/SENuserdocs/SENuser.html#overview
- P. Gary, et al, "New 10-Gbps Networks Facilitating Grid-related Development Activities," NASA GSFC Grid Workshop, Lanham, MD, June, 2005, especially pp. 26-31: http://romulus.gsfc.nasa.gov/msst/gridws/P19%20Gary.pdf
- L-Net Project website: http://cisto.gsfc.nasa.gov/IRAD_Lambda.html

Network Path and Application Diagnosis

(NPAD/*Pathdiag)*

*(*http://www.psc.edu/networking/projects/pathdiag/)

Recommended NCCS+Columbia User "Script"

- From your host's web-browser link to above URL and then to any of the NPAD servers identified therein
  - Best if nearest, so for now use: http://kirana.psc.edu/NPAD/
- In NPAD server's data entry box, enter:
  - "88" for Round Trip Time (msec)
  - "5" for Target Rate (Mbps)
    - Also try "50" and "90"
- After testing, email sen-help@sci.gsfc.nasa.gov citing your hostname and which NPAD server you used

Some Potentially Applicable "Futures"

"Under Development" Presented at ESCC/Internet2 Joint Techs Workshop, Madison, July, 2006

- S. Shalunov, "VFER: developing a bulk transport tool within Google Summer of Code": http://events.internet2.edu/2006/jt-madison/sessionDetails.cfm?session=2782&event=253

- G. Almes et al, "Phoebus: Achieving Dependable Bulk Throughput in a Hybrid Network": http://events.internet2.edu/2006/jt-madison/sessionDetails.cfm?session=2754&event=253

- E. Boyd et al, "Performance Update (& perfSONAR)": http://events.internet2.edu/2006/jt-madison/sessionDetails.cfm?session=2799&event=253


NETWORK
BOTTLENECKS

GSFC SEN and HECN

SEN+HECN User Help Desk

- SEN user registration
- Firewall waiver requests
- Throughput performance tuning
- Troubleshooting
- SEN Alert/Notices
- Available 24/7 ad hoc
  - sen-help@sci.gsfc.nasa.gov
  - 301-286-2045
- http://cisto.gsfc.nasa.gov/SENuserdocs/SENuser.html

Meeting for Network-Affected NCCS+Columbia Users

Thursday, 14Sep06, 1-2pm, in GSFC building 33, room H114

Backup Slides

Internet2 Land Speed Record

(Rules and current records: http://lsr.internet2.edu/)

Near-to-

Last IPv4 Single Stream Record (http://data-reservoir.adm.s.u-tokyo.ac.jp/lsr-20041225/)

- **7.21 Gbps** (TCP payload), standard frame, 216,300 petabit-meters / second
- 33,979 km connection from Tokyo through Chicago, Amsterdam, NewYork to Tokyo Latency 499ms RTT

NETWORK USED IN THE EXPERIMENT: nearly three-quarters the circumference of the Earth

The Wizard Gap

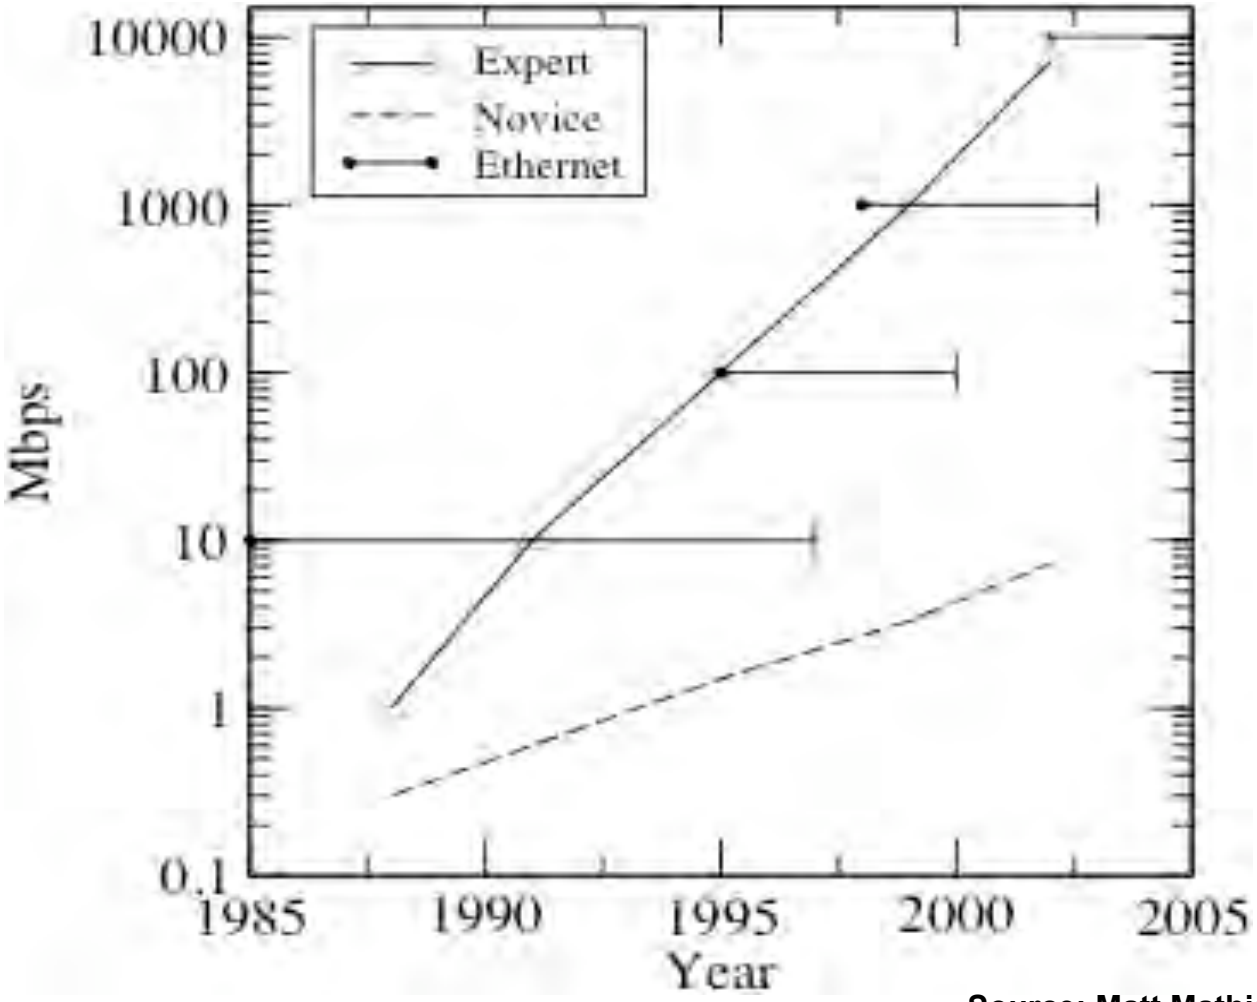

Source: Matt Mathis (PSC)

The Wizard Gap Updated

- Experts have topped out end systems & links
  - 10 Gb/s NIC bottleneck
  - 40 Gb/s “link” bandwidth (striped)
- Median I2 bulk rate is 24 Mbit/s
  - 3 M Byte/s
  - See http://netflow.internet2.edu/weekly/
- Current Gap is about 1000:1
- Closing the first factor of 10 should now be “easy”

Source: Matt Mathis (PSC)

NPAD/Pathdiag - Why should you care?

- One click automatic performance diagnosis
  - Designed for (non-expert) end users
  - Accurate end-systems and last mile diagnosis
    - Eliminate most false pass results
    - Accurate distinction between host and path flaws
    - Accurate and specific identification of most flaws
  - Basic networking tutorial info
    - Help the end user understand the problem
    - Help train 1st tier support (sysadmin or netadmin)
    - Backup documentation for support escalation
- Empower the user to get it fixed
  - The same reports for users and admins

Source: Matt Mathis (PSC)

Recalibrate user expectations

- Long history of very poor network performance
  - Users do not know what to expect
  - Users have become completely numb
- Goal for new baseline user expectations:
  - 1 Gigabyte in less than 2 minutes (~67 Mb/s)
- Everyone should be able to reach these rates by default
- People who can't should know why or be very angry

Source: Matt Mathis (PSC)

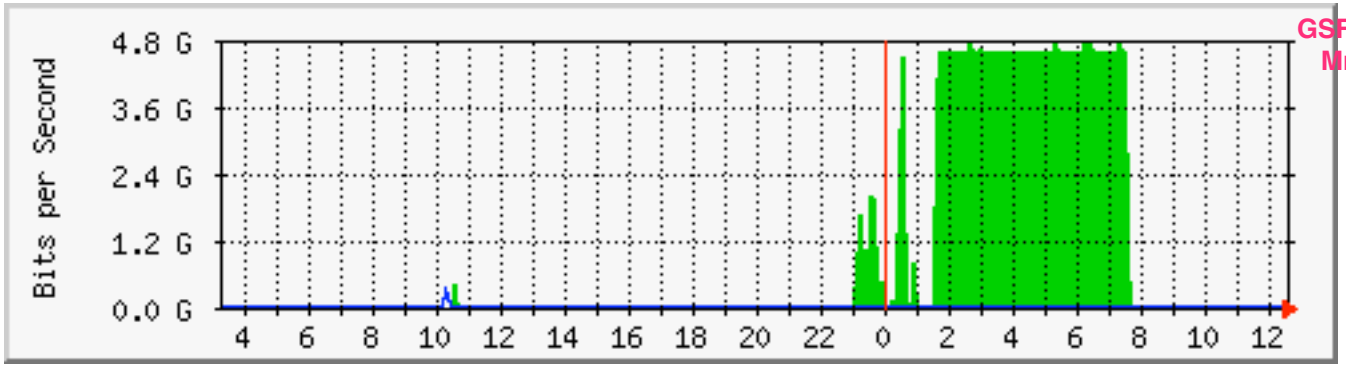

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

chance4 10 GigE

3 February 2006

Max In:	4772.6 Mb/s (47.7%)	Average In:	909.6 Mb/s (9.1%)	Current In:	0.0 b/s (0.0%)
Max Out:	379.1 Mb/s (3.8%)	Average Out:	1826.5 kb/s (0.0%)	Current Out:	8.0 b/s (0.0%)

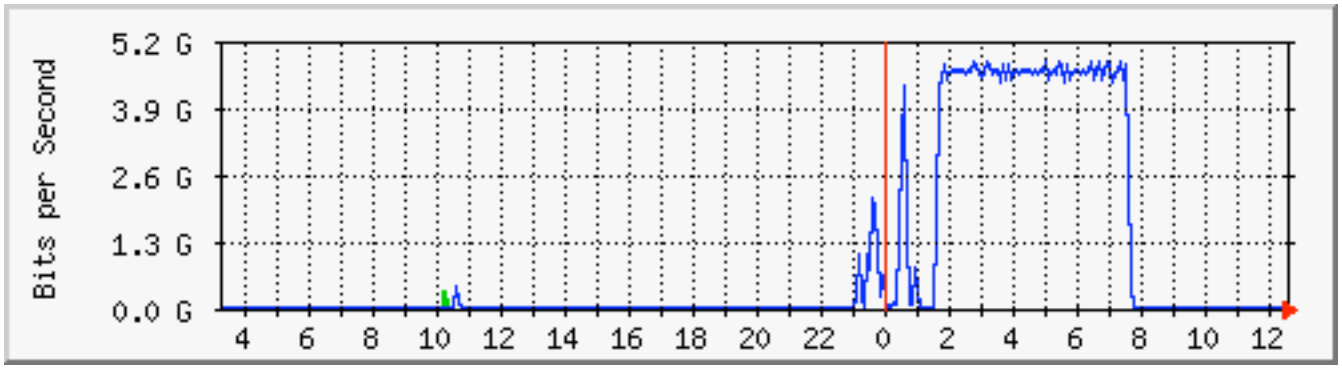

chance5 10 GigE

3 February 2006

Max In:	386.7 Mb/s (3.9%)	Average In:	1815.2 kb/s (0.0%)	Current In:	0.0 b/s (0.0%)
Max Out:	4817.9 Mb/s (48.2%)	Average Out:	907.1 Mb/s (9.1%)	Current Out:	8.0 b/s (0.0%)

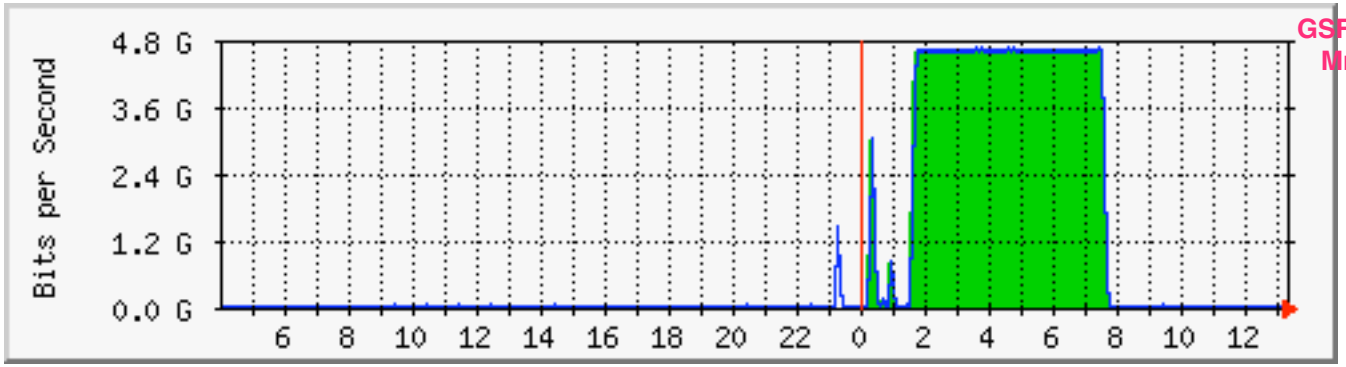

GSFC Scientific and Engineering Network (SEN)
Mrtg-based `Daily' Graph (5 Minute Average)
Bits per second In and Out
On Selected Interfaces

GSFC-McLean 10 GigE "Lambda49" across DRAGON

3 February 2006

Max In:	4690.0 Mb/s (46.9%)	Average In:	857.2 Mb/s (8.6%)	Current In:	552.0 b/s (0.0%)
Max Out:	4687.6 Mb/s (46.9%)	Average Out:	861.4 Mb/s (8.6%)	Current Out:	96.0 b/s (0.0%)

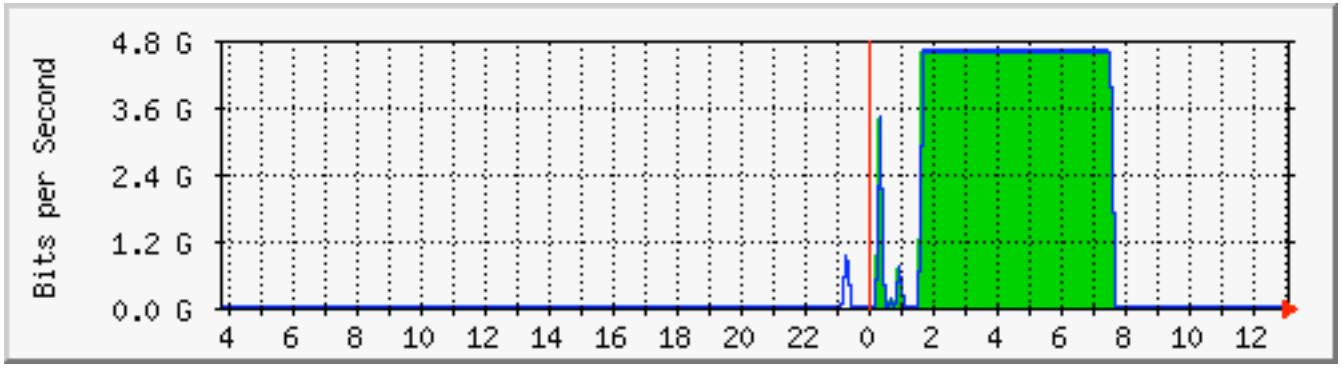

Both chance and chance2 via RAPTOR switch at GSFC 10 GigE

3 February 2006

Max In:	4653.9 Mb/s (46.5%)	Average In:	852.4 Mb/s (8.5%)	Current In:	0.0 b/s (0.0%)
Max Out:	4651.7 Mb/s (46.5%)	Average Out:	856.7 Mb/s (8.6%)	Current Out:	400.0 b/s (0.0%)

High Level View of GSFC SEN Links with Major User Organizations (pre GSFC SED Transformation)

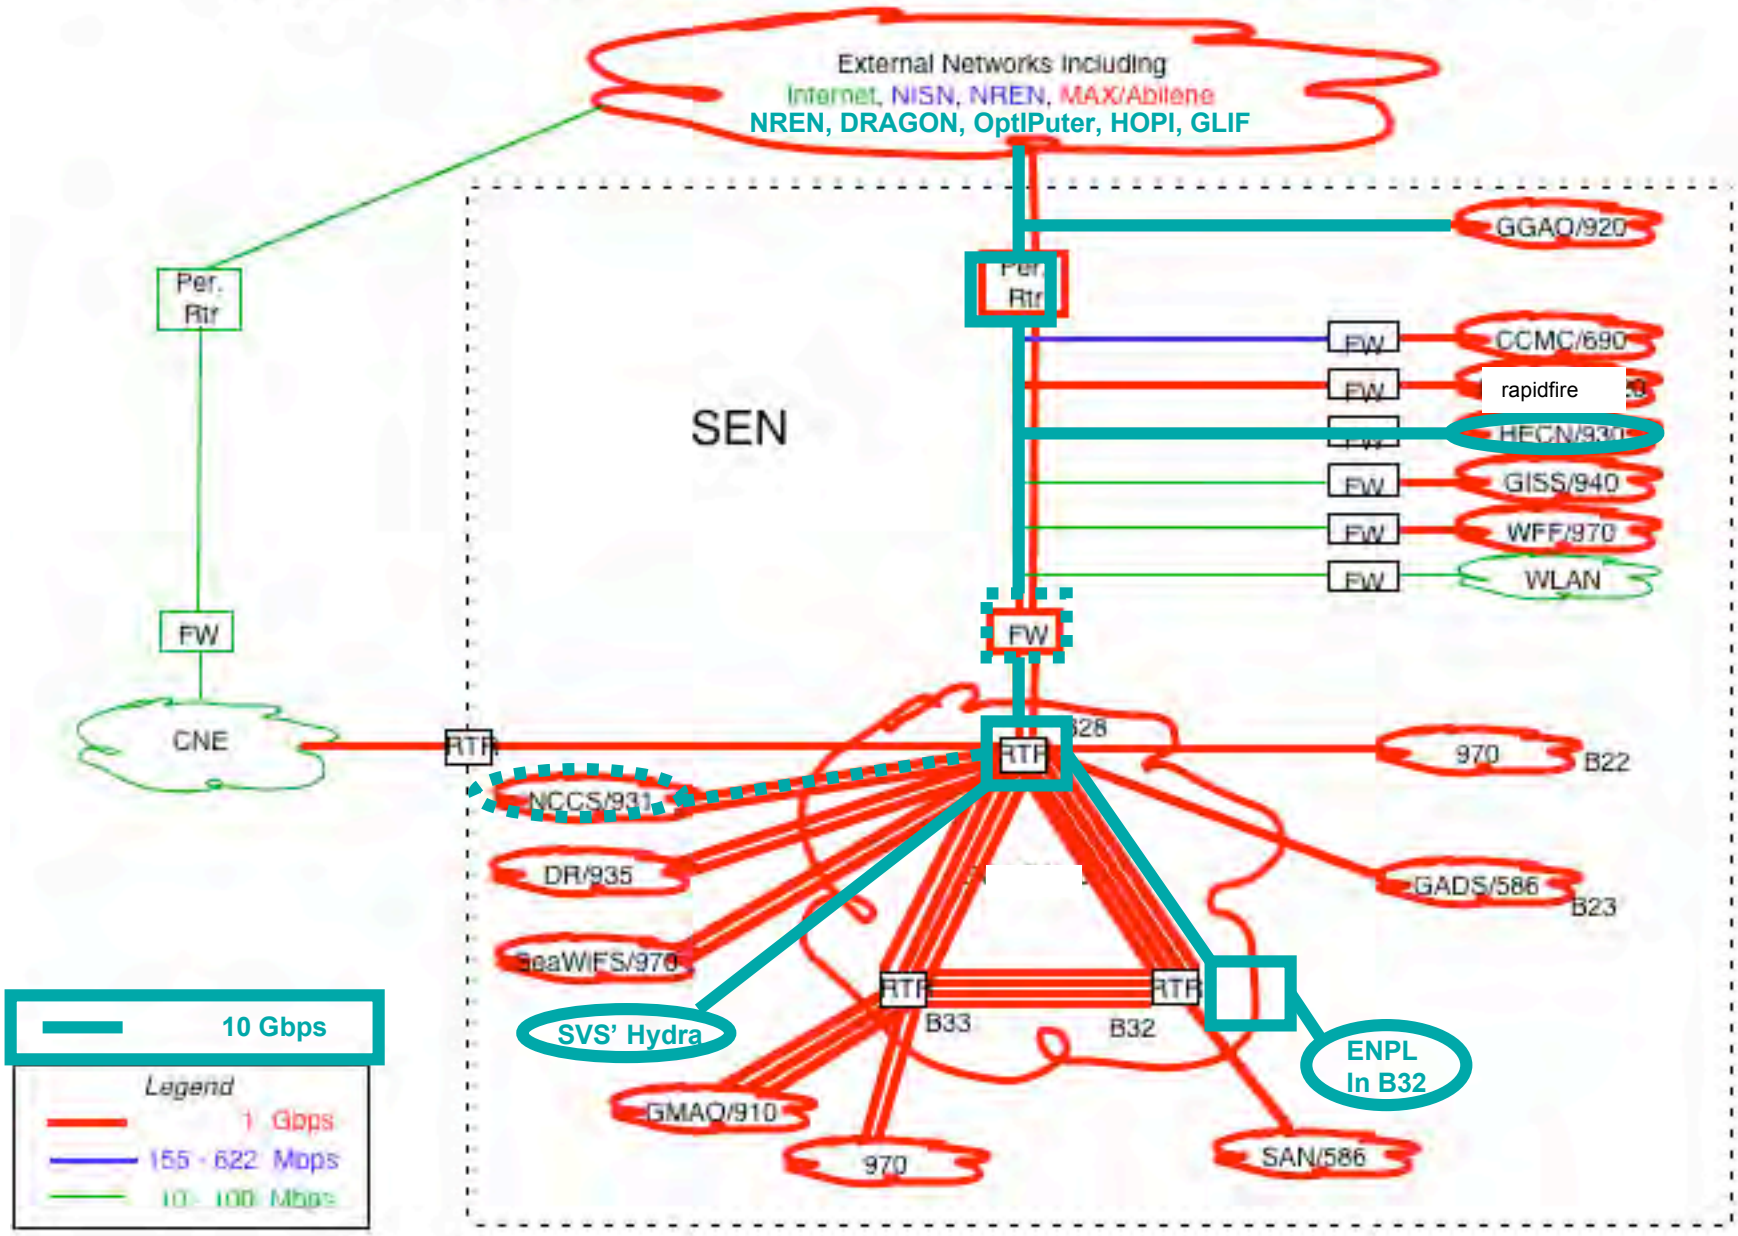

02/23/06

JPG 01/19/05

GSFC's Scientific and Engineering Network (SEN) Overview

Some Key Data Flow/Path Restrictions Affecting CNE & SEN

[By GSFC security policy and/or management agreement]

- Only CNE-sourced traffic with SEN-destinations
  or SEN-sourced traffic with CNE-destinations
  can flow across the "Only intra-CNE-SEN routing" path between CNE & SEN

- CNE-source traffic can **not** traverse the SEN or its Border Router/FW to other networks; SEN-source traffic can **not** traverse the CNE or its Border Router/FW to other networks
  - Rationale:
    - The security policies/rulesets of the respective Border Routers/FWs intentionally block data flows transiting to/from other networks
    - The respective CNE & SEN Helpdesks are not adequately trained or equipped to support data flows transiting to/from other networks

The Only Flows Permitted Across Intra-CNE-SEN Router By GSFC Security Policies and/or Management Agreements

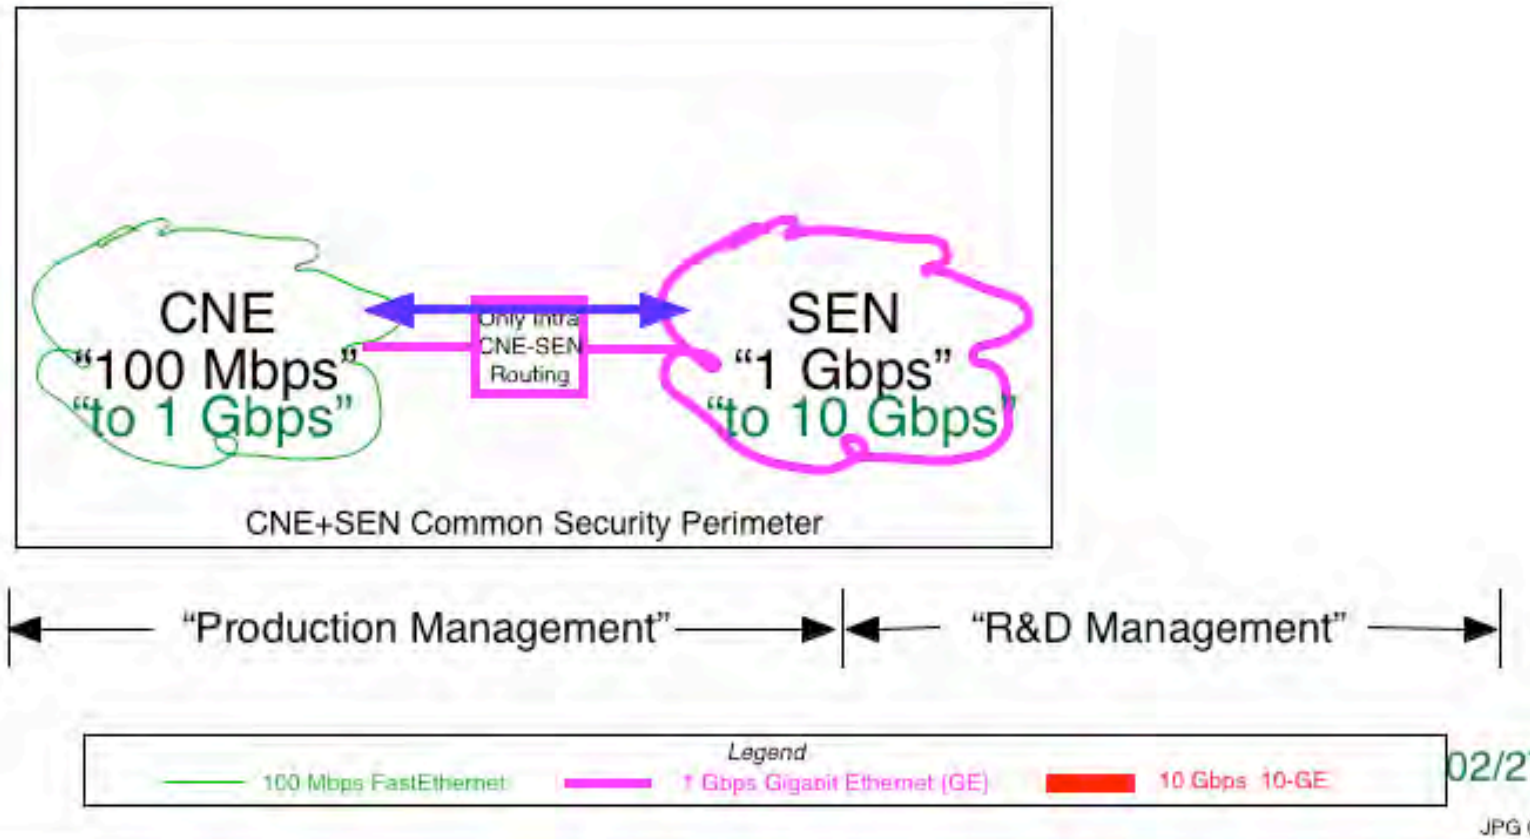

Examples of Traversal Flows Not Permitted By GSFC Security Policies and/or Management Agreements

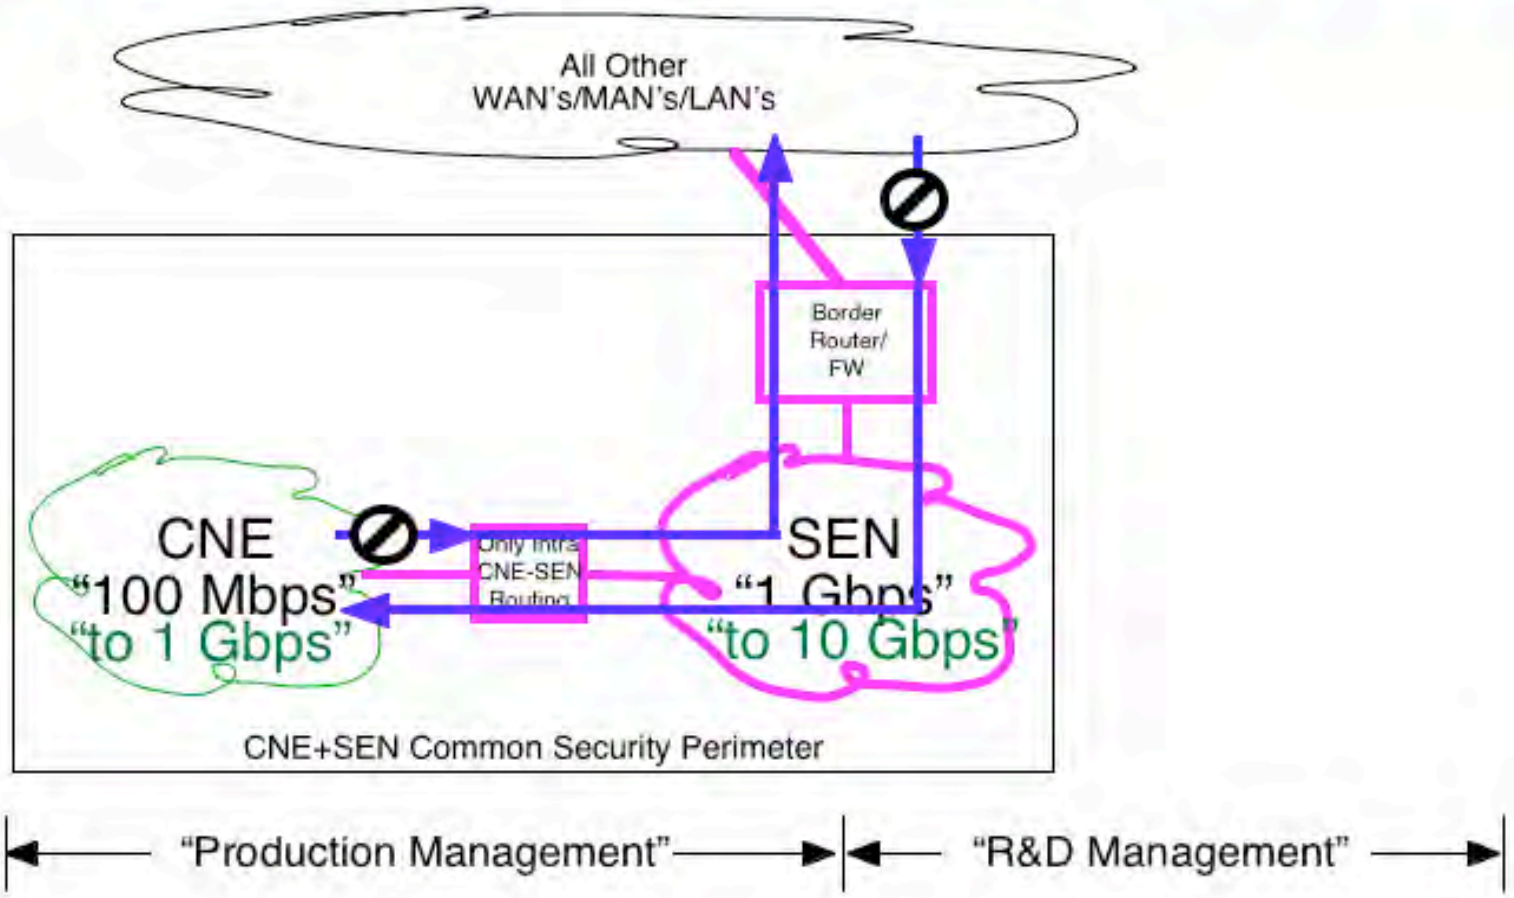

02/27/06

JPG 02/25/04

GSFC SEN and HECN Local Connectivity/Access

Wired

- Inter-building cable plant
  - 10-Gbps: B28+B32, B28+B201(GGAO)
  - 4-Gbps (via IEEE 802.3ad link aggregation standards): Among B28, B32, B33
  - 1-Gbps: B28+B22, B28+B23
  - 30-45 Mbps: GISS (NYC), WFF
- Intra-building connections with local user desktops/servers/clusters
  - Minimum 1-GE (~280 at Greenbelt, ~200 at GISS, ~125 at WFF)
  - 10-GE (~20 at Greenbelt)

GSFC SEN and HECN Local Connectivity/Access

Wired (continued)

- 9000 byte Ethernet jumbo frames supported throughout entire local infrastructure

WLAN

- 802.11 b & g (~100 for Code 606 and their guests in B28 only)
- Security protections
  - Rules of Behavior signed by users
  - MAC addresses registered and used for authentication
  - NAT to private address space
  - WPA2 encryption
  - SSID not to be broadcast (future, but soon)

GSFC SEN and HECN Remote Access

- "Production" WAN's via separate connections
  - 10-GE with NREN (http://www.nren.nasa.gov/)
  - 2-GE with Mid-Atlantic Crossroads (MAX) regional gigapop (http://www.maxgigapop.net/index.html) and Internet2's Abilene network (http://abilene.internet2.edu/)
  - 1-GE with GSFC isolan switch for connecting with CNE, GSFC-based mission networks, and NISN
- "R&D" WAN's via multi-wavelength division multiplexing
  - 2.4- & 10-Gbps with UMCP-led DRAGON regional optical network (http://dragon.maxgigapop.net/)
  - 10-Gbps with UCSD-led OptIPuter Project (http://www.optiputer.net/) via the National LambdaRail (NLR) (http://www.nlr.net/)
  - 10-Gbps with NGC-led BICC lab (http://www.it.northropgrumman.com/index.html)

GSFC High End Computer Network (HECN) Project's Research Partners and Collaborators

- *DRAGON Project: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/WebHome*
  - *PI: Jerry Sobieski (UMCP)*
  - *GSFC L-Net on DRAGON network diagram: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/Network*
- *e-VLBI Project: http://web.haystack.mit.edu/e-vlbi/evlbi.html*
  - *PI: Alan Whitney (MIT/Haystack)*
  - *GSFC L-Net on e-VLBI network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SC04_eVLBI_network.pdf*
- *GLIF: http://www.glif.is/*
  - *Chair: Kees Neggers (SURFnet)*
  - *GLIF network diagrams: http://www.glif.is/publications/#maps*
- *NGC IT Sector: http://www.it.northropgrumman.com/index.html*
  - *PI: Brice Womack (NGC)*
  - *GSFC L-Net on NGC IT Sector Colshire network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/DRAGON_NGC_030606.pdf*
- *NLR: http://www.nlr.net/*
  - *CEO: Tom West (NLR)*
  - *NLR network diagram: http://www.nlr.net/infrastructure/*
- *NREN Project: http://www.nren.nasa.gov/*
  - *PM: Ken Freeman (ARC)*
  - *GSFC L-Net/SEN on NREN network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/CENIC2006_13_mfoster_excerpts.pdf*
- *OptIPuter Project: http://www.optiputer.net/*
  - *PI: Larry Smarr (UCSD)*
  - *GSFC L-Net on OptIPuter network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SMARR-OptIPuter-AHM-gold.pdf*
- *TeraFlow Testbed Project: http://www.teraflowtestbed.net/*
  - *PI: Robert Grossman (UIC)*
  - *GSFC L-Net on TeraFlow Testbed network diagram: http://www.ncdm.uic.edu/maps/index.jpeg*

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- eVLBI (w/MIT-Haystack, …)
- OptIPuter & Multi-channel Collaboration/Video Streaming Technologies(w/UCSD & UIC)
- 3D HDTV-over-IP R&D (w/Physical Optics Corporation)
- Distributed ESMF R&D
- SAN-over-IP (w/UMIACS & NGC)
- Using ARC/NAS/Columbia Supercomputer (w/NREN)

DRAGON eVLBI Experiment Configuration

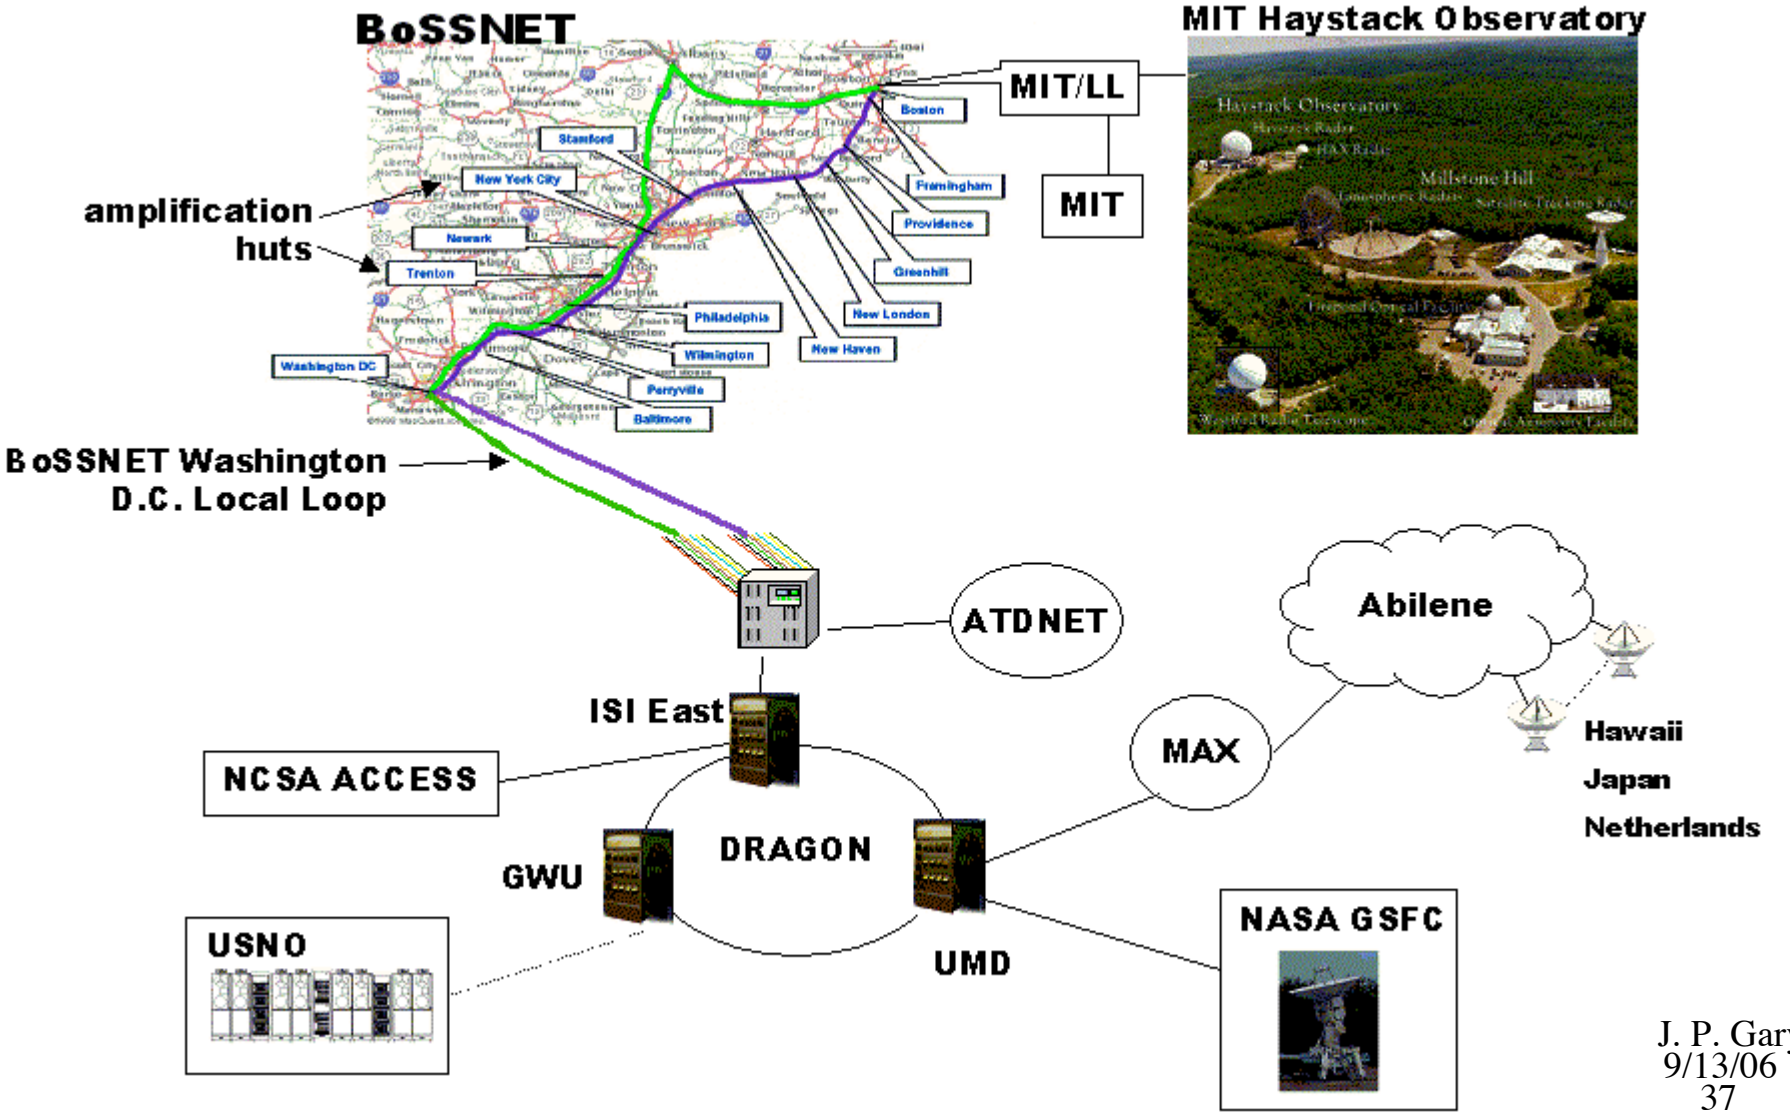

NASA GSFC Tests with OptIPuter Across the National LambdaRail

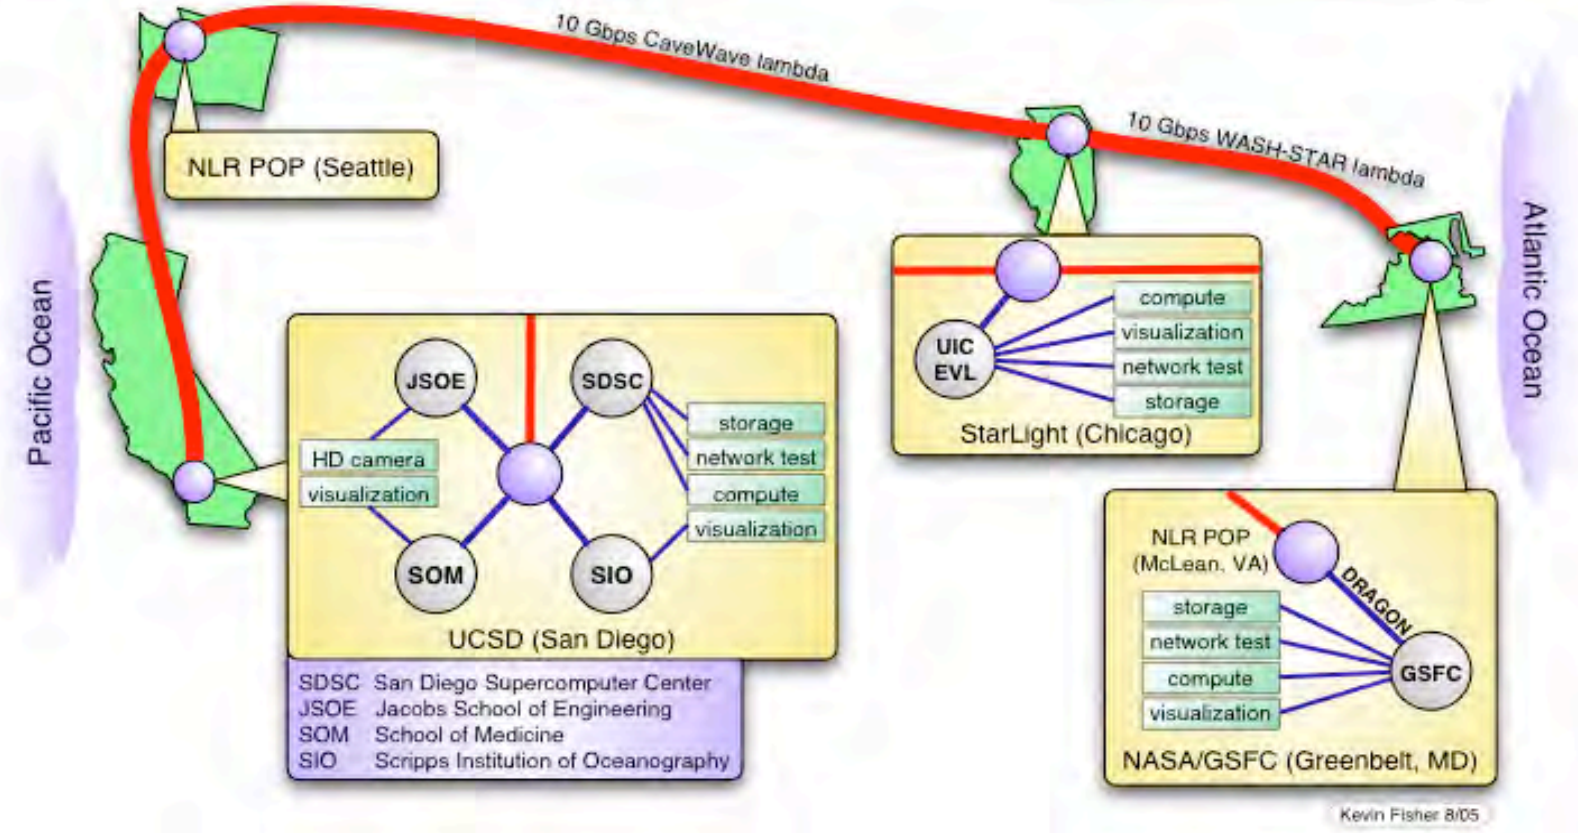

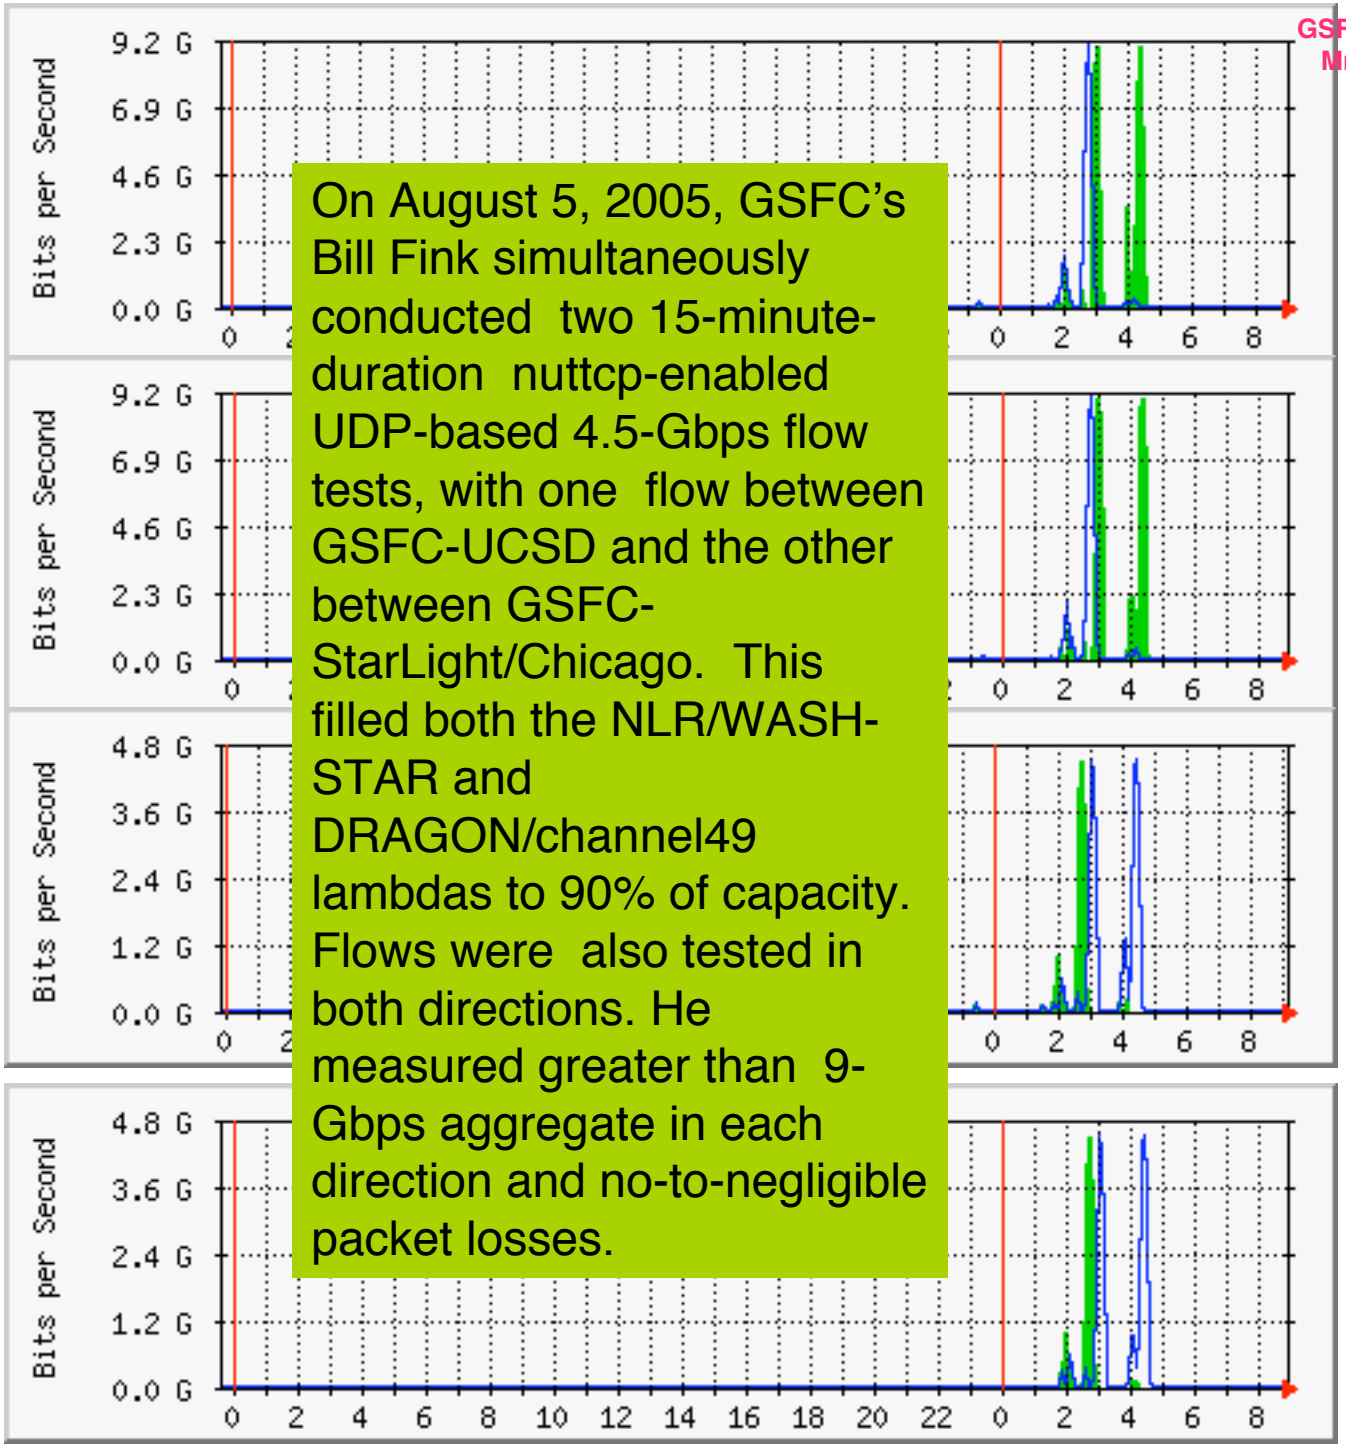

On August 5, 2005, GSFC's Bill Fink simultaneously conducted two 15-minute-duration nuttcp-enabled UDP-based 4.5-Gbps flow tests, with one flow between GSFC-UCSD and the other between GSFC-StarLight/Chicago. This filled both the NLR/WASH-STAR and DRAGON/channel49 lambdas to 90% of capacity. Flows were also tested in both directions. He measured greater than 9-Gbps aggregate in each direction and no-to-negligible packet losses.

10 GigE from McLean to Chicago via OptIPuter Lambda
5 August 2005

DRAGON 10Gig DWDM XFP
5 August 2005

chance1 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

chance2 10Gig (eth1 Intel Pro/10GbE)
5 August 2005

Combining Telepresence with Remote Interactive Analysis of Data Over NLR

www.calit2.net/articles/article.php?id=660

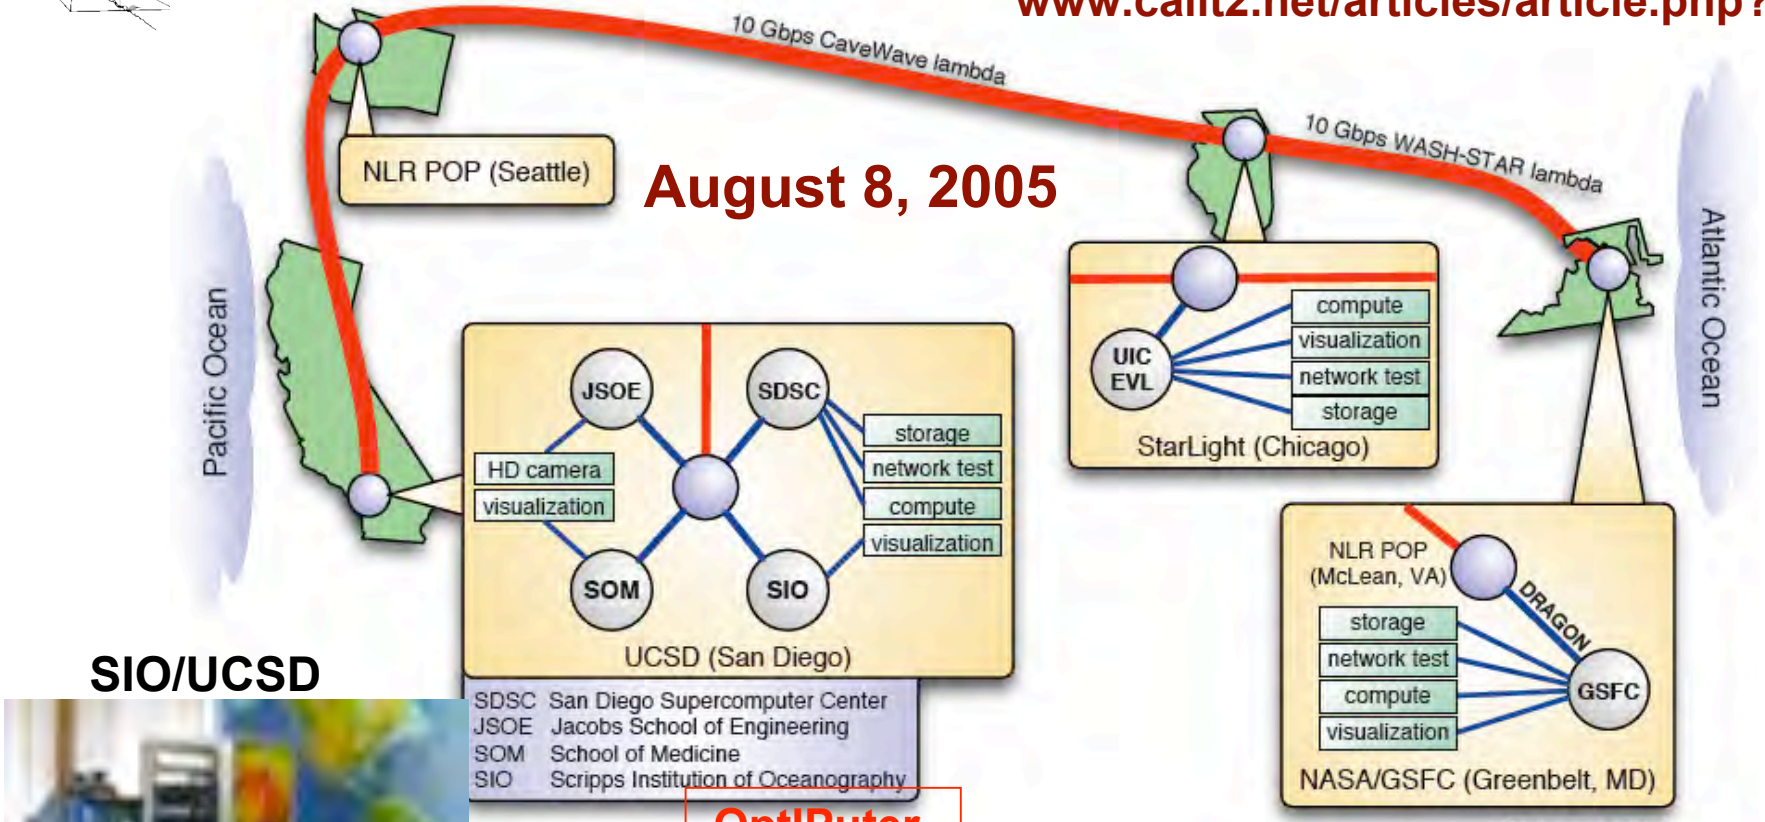

SIO/UCSD

OptIPuter Visualized Data

HDTV Over Lambda

NASA Goddard

Gold Standard OptIPuter Initial Sites
Persistent 10G Services with SAGE End Appliances

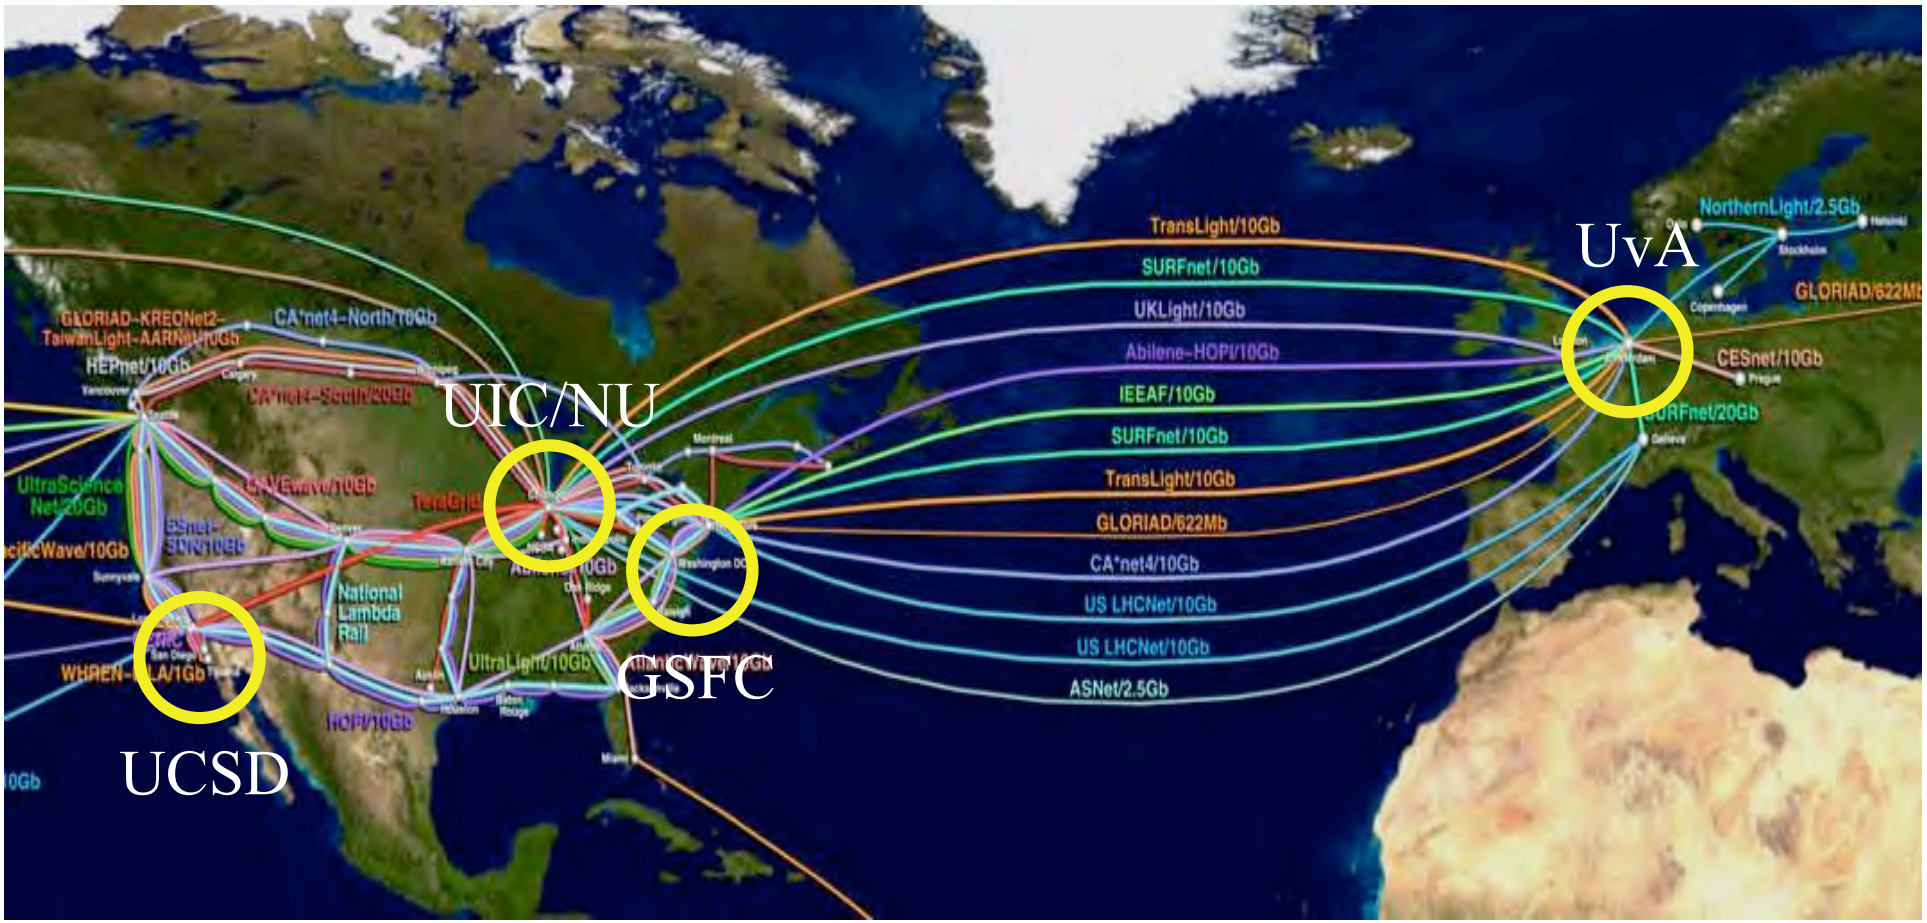

OptIPuter: Empowering Global Data Intensive Services for Science

- Observations
  - The OptIPuter Project Has Moved Beyond “Proof of Concept”
  - Although the OptIPuter is a Technical Achievement, the Research Group Defines Success by Application Achievements
- Goals
  - Provide Persistent, Reliable Unique Services for Data Intensive Sciences and Applications
  - Design, Develop and Deploy Gold Standard OptIPuter Implementation, Which Will Provide Those Services
  - The “Gold Standard” OptIPuter Will Comprise a Distributed Contiguous, Foundation Facility
  - Key Focal Services – e.g., SAGE Distributed via Rocks
  - UltraHD Personal Collaborative Communications Directly Integrated With Extremely Large-Scale, Real-Time, Data Streaming
  - Flexible, Scalable OptIPuter Environments

iGrid 2005 Workshop, 26-29Sep05, UCSD/CalIT2

Accelerating the Use of Multi-10Gigabit per Second International and National Networks: www.igrid2005.org

GSFC's Ben Kobler (left) and POC's Sookwang Ro and Kirill Kolesnikov (right) work to set up POC's 35" x 35" holographic 3D HDTV video display system (center) prior to the start of iGrid 2005.

US130: Real-Time True-3D/HDTV (No Goggles) Visualization Over the National LambdaRail

NASA and Physical Optics Corporation demonstrate a holographic 3D HDTV video display system that does not require goggles or other special head gear, using a live cross-country video feed from NASA Goddard Space Flight Center to the iGrid 2005 site in San Diego. POC is a NASA SBIR Phase 1 awardee, and worked with NASA GSFC on this project.

www.poc.com/emerging_products/3d_display/default.asp

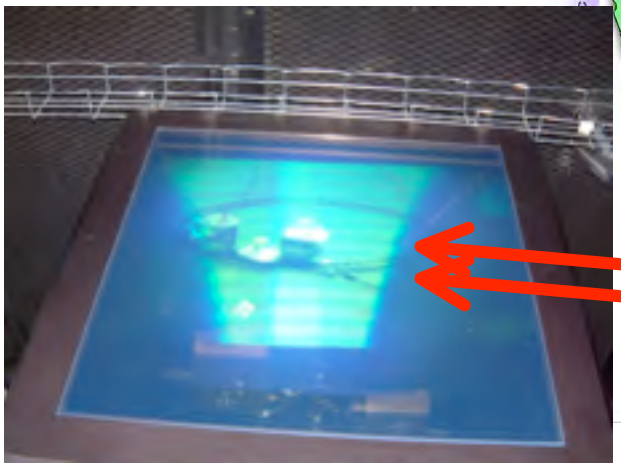

Only a non-stereo image of the True-3D display is captured in this photo of the real-time stereo-HDTV images transmitted from GSFC.

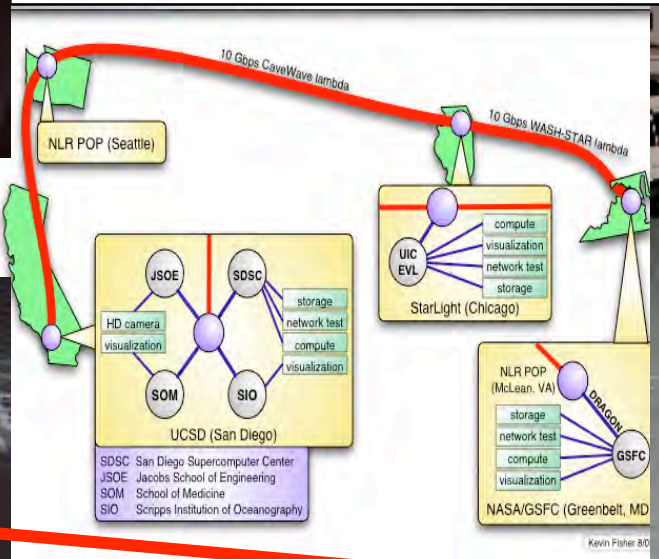

3D HDTV Over Lambda

Stereoscoptically-aligned Sony HDV 1080i HDR-FX1HDTV cameras and the viewed targets at GSFC.

Cross-Organization Coupling of Climate Models through ESMF

(A Prototype Over High-Speed Networks)

Shujia Zhou (Lead), C. Cruz, R. Burns, B. Womack, G. Higgins NASA SIVO/Northrop Grumman TASC

Collaborators:

- **High-speed network: P. Gary, B. Fink, P. Lang (NASA GSFC/ADNET)**
- **Cluster system admin: K. Fisher (NASA GSFC)**
- **XCAT/Proteus: M. Govindaraju, K. Chiu, M. Head (SUNY, Binghampton)**
- **Models: J. Spahr, C. Mechoso (UCLA), C. Hill (MIT), P. Jones (LANL)**

Presented at NASA Exhibit (booth 1810) at SC|05, November 14-18, 2005

EARTH SYSTEM MODELING FRAMEWORK

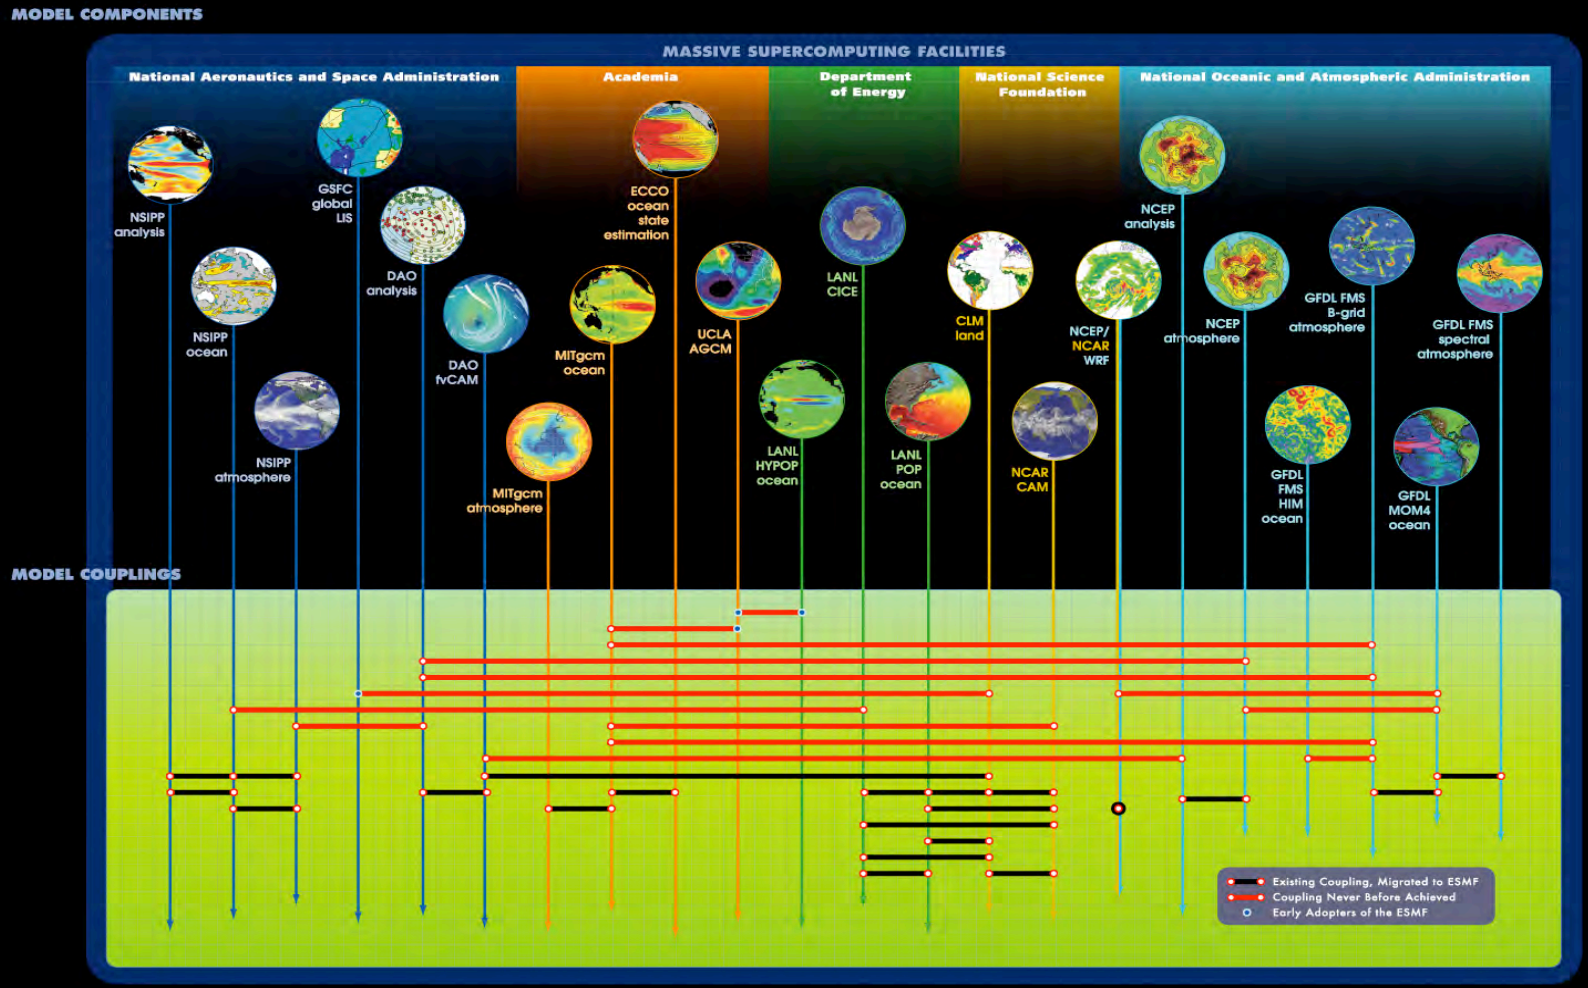

Coupled Atmosphere-Ocean Models

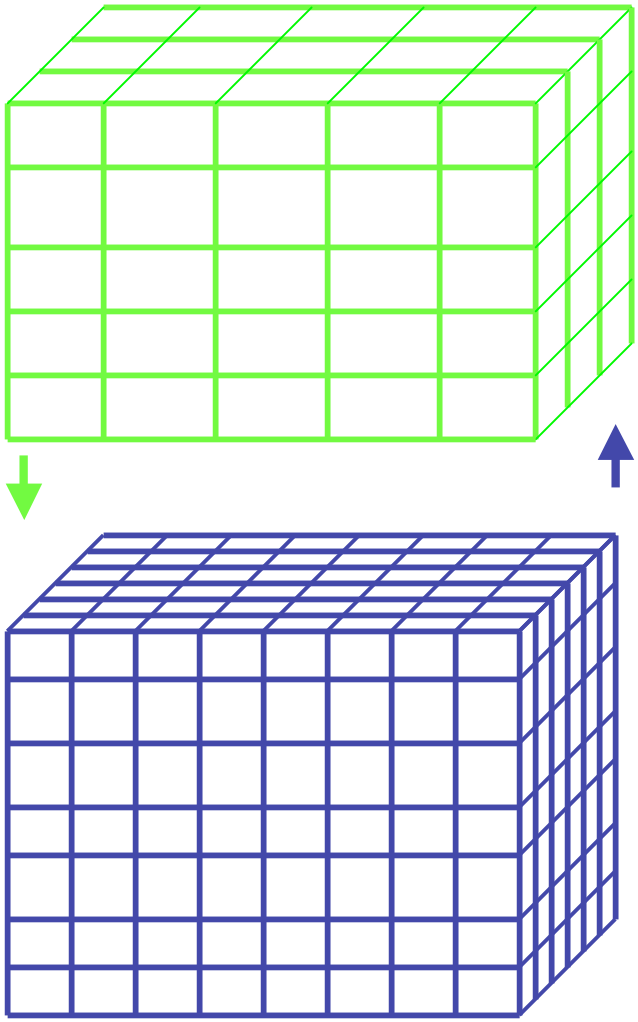

Atmosphere

Ocean

Different grid type, resolution

Thunder2-Thunder3 Looped Data Flows Across DRAGON Regional Optical Network

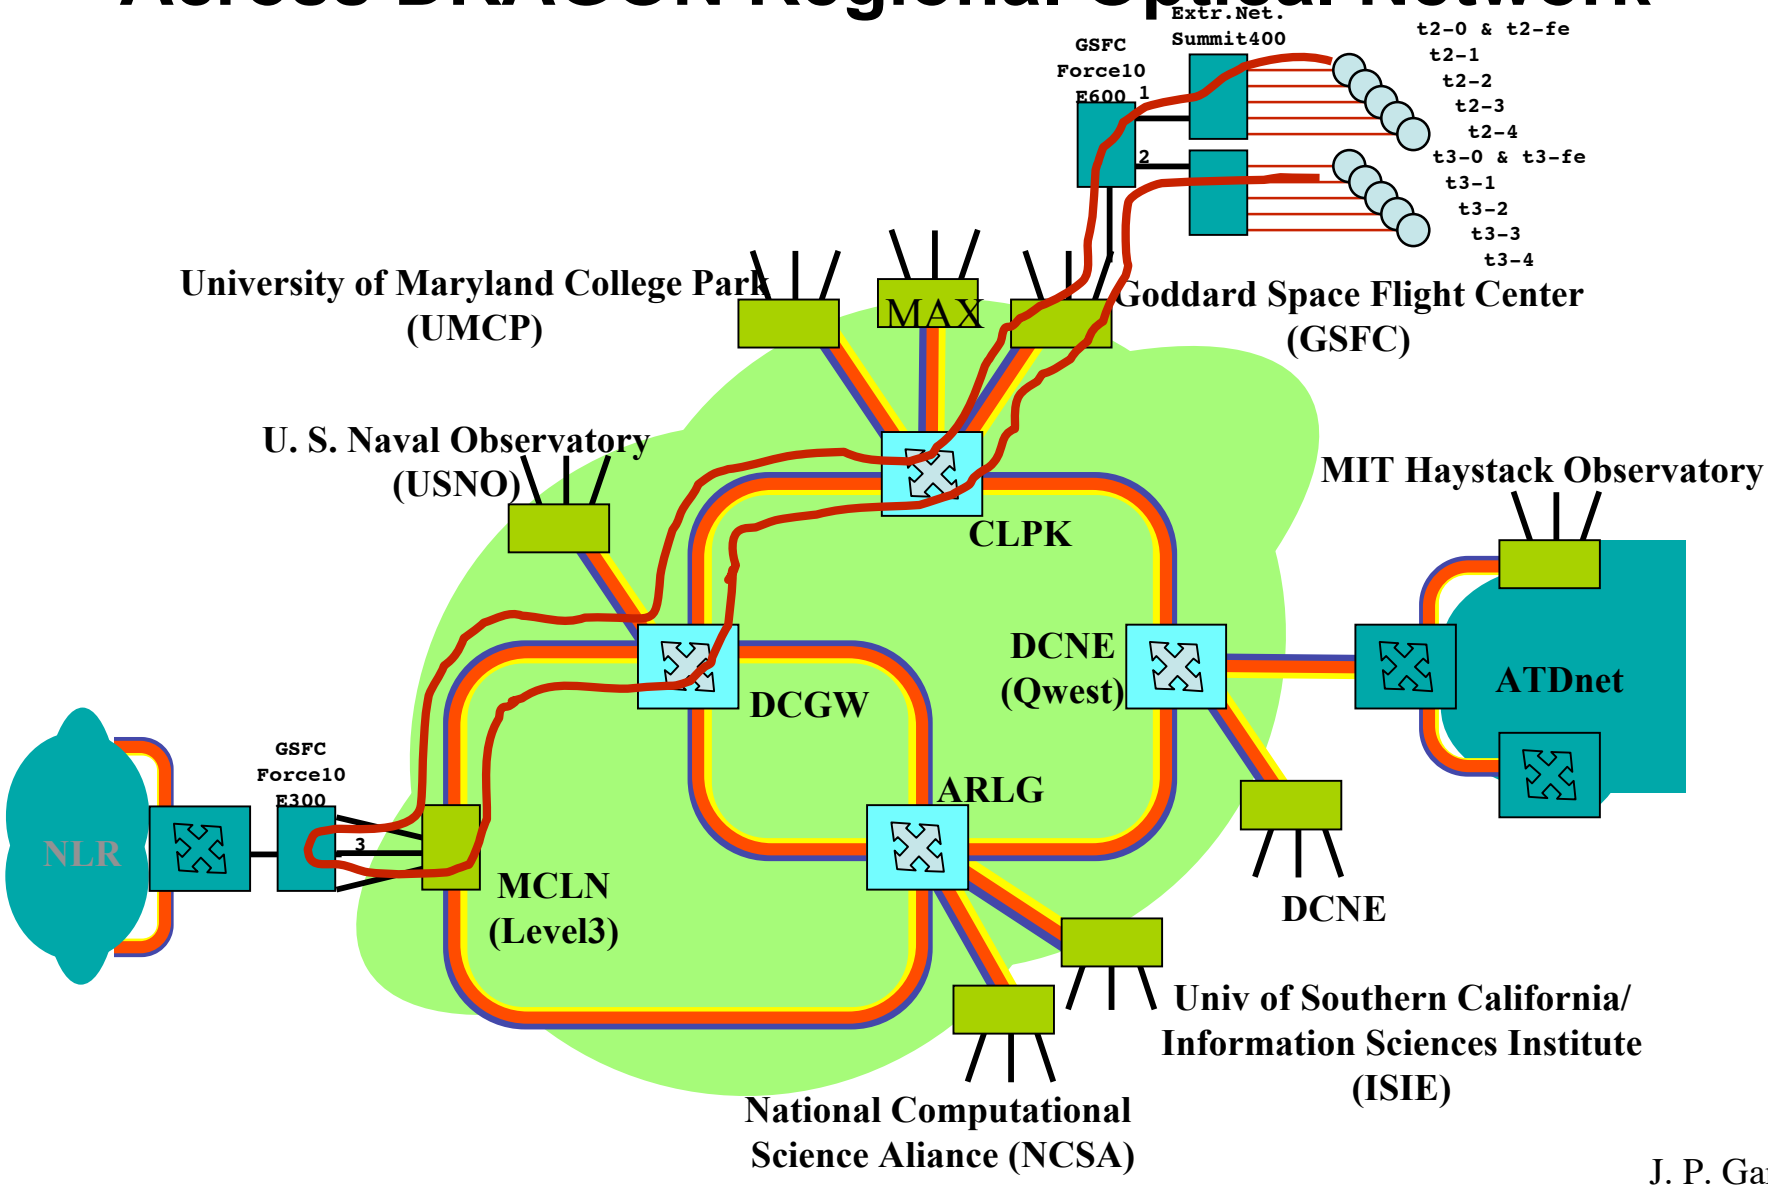

Current SAN-over-IP Test-bed

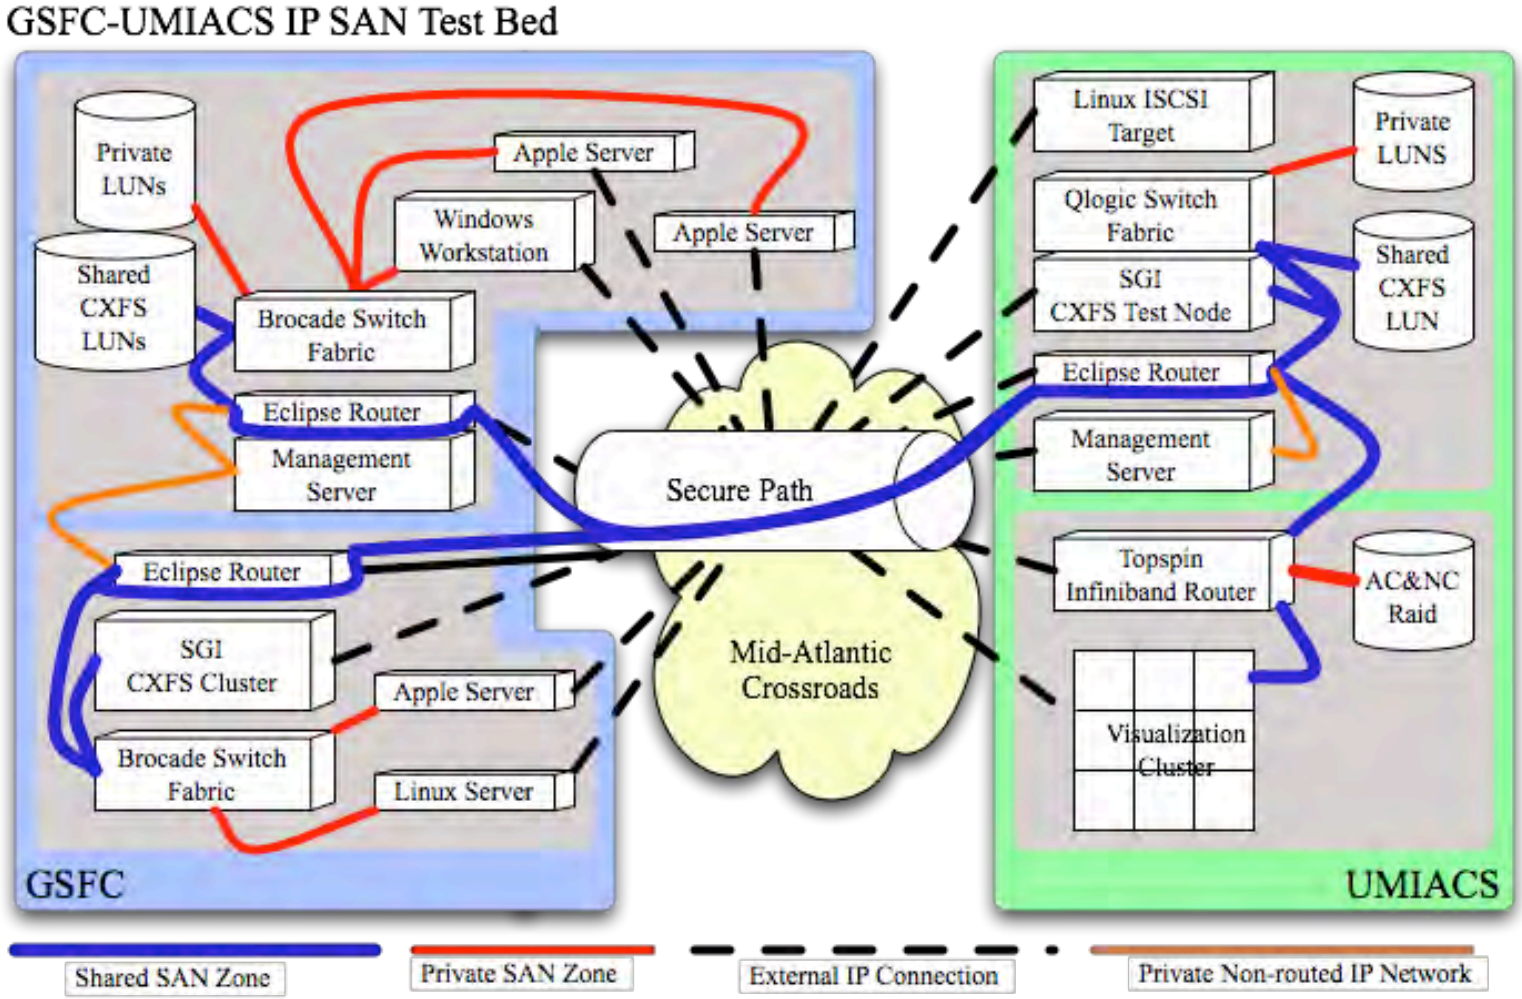

Source: Fritz McCall (UMIACS)

Wide Area Storage Configuration

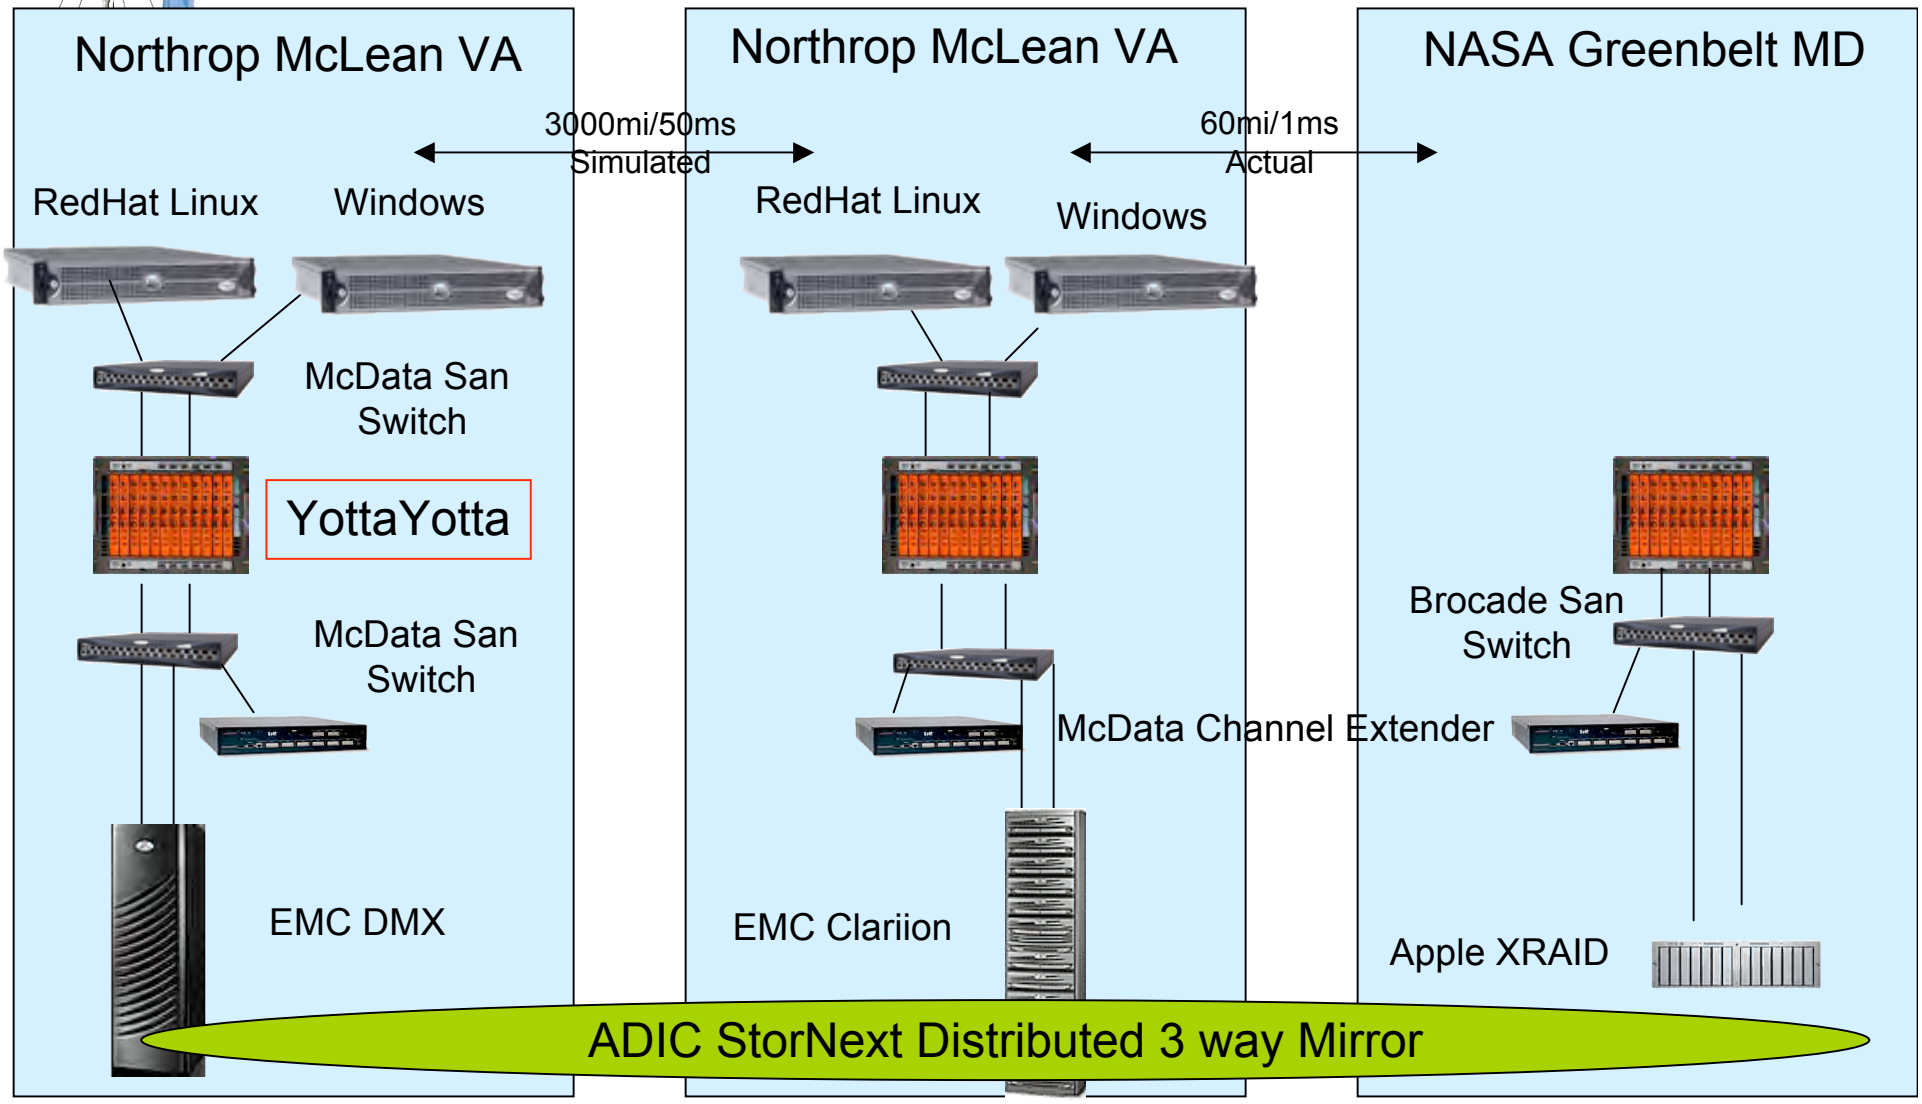

Source: Bob Bramow (YottaYotta)

Columbia Supercomputer

- 10,240 1.6 GHz CPUs
- Configured as twenty 512 CPU single-system image nodes via NUMA
- SGI Altix 3700 Architecture, runs Linux
- 1 Terabyte shared memory per node
- Over 500 terabytes of online disk space

WAN Infrastructure: NREN

Target CY06

10G waves at the core, dark fiber to end sites

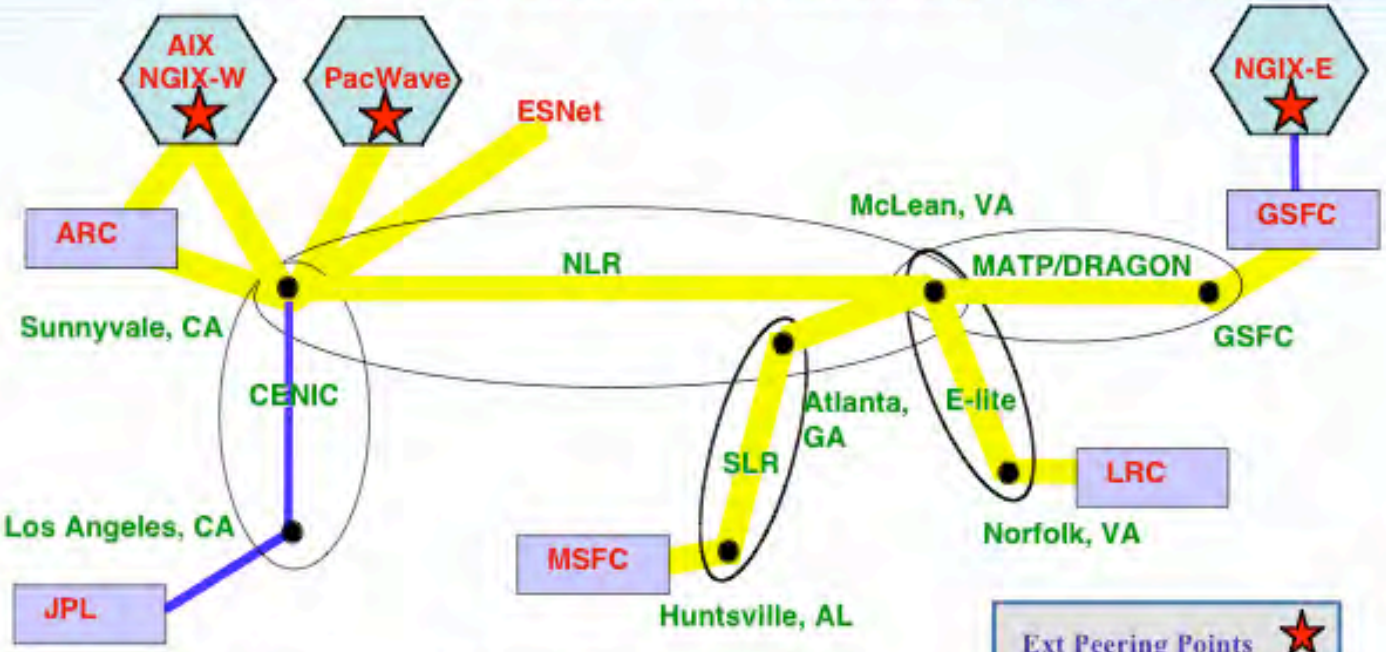

- National and Regional optical networks provide links over which 10 Gbps and 1 Gbps waves can be established.
- Distributed exchange points provide interconnect in metro and regional areas to other networks and research facilities

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- Multi-channel Collaboration/Video Streaming Technologies
  - Scalable Adaptive Graphics Environment (SAGE) (http://www.evl.uic.edu/cavern/sage)
  - HDTV-over-IP
  - Demonstrations of 21st Century National-Scale Team Science (http://www.calit2.net/newsroom/release.php?id=660)
- 3D HDTV-over-IP
  - 3D Multichannel Networked System via NASA SBIR FY06 Phase2 awardee Physical Optics Corporation
  - Live 3D HDTV multi-Gbps real-time data streaming from GSFC to holographic display at iGrid2005 as the US130/Real-Time_True-3D_Visualization exhibitor (http://www.igrid2005.org/program/applications/vizservices_3dviz.html)

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Previous and/or On-Going Applications

- Enabling e-VLBI real-time data flows from GGAO to MIT/Haystack (http://web.haystack.mit.edu/e-vlbi/evlbi.html)
- Prototyping of Earth System Modeling Framework (ESMF)-based cross-organization coupling of climate models over a high speed network (http://cisto.gsfc.nasa.gov/L-Netpdfs/sc05_esmf_demo_v5.pdf)
- Evaluating SAN-over-IP and distributed shared file system applicability to enhancing science data flows
  - NCCS' participation in the Data Intensive Computing Environment (DICE) Project (http://www.avetec.org/dice)
  - NCCS data portal environment
  - McCall et al, "A framework for Managing Inter-site Storage Area Networks using Grid Technologies" (http://romulus.gsfc.nasa.gov/msst/conf2006/Papers/2006-025-McCall.pdf)

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Upcoming and/or Future Applications

- MAP'06 (w/NGC)
- Phoebus (w/Internet2, UDel & GSFC/ENPL)
- SOA+Brokering for ECHO (w/SIO, JPL & UAH)
- Dynamic Linking (w/ORNL, CUNY)
- Grid Computing (w/TBD: SURAGrid, UMBC, …)

2006 Hurricane Season - Global Modeling

NCEP
Inputs (1 GB)

INTEL Corp
Portland, Oregon

NASA Ames
Mt. View, California

Northrop Grumman
McLean, Virginia

NASA Goddard
Greenbelt, Maryland

DISTRIBUTED COMPUTING NODES

NEXT-GEN NETWORKS

Conventional Network
(600 Mb/s)

National Lambda Rail (10 - 40 Gb/s)

All model Outputs (500 GB per execution, ~15TB for entire season

Tape Backup

NASA Goddard
Greenbelt, Maryland

Main Server

NASA Goddard
Greenbelt, Maryland

DATA SERVERS / LONG TERM STORAGE

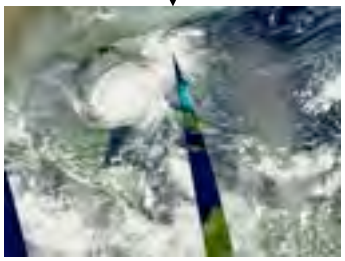

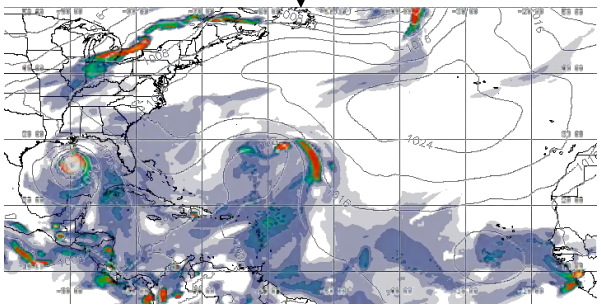

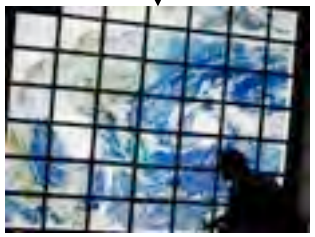

WEB SERVICES / ADVANCED VISUALIZATIONS

Source: Mike Seablom (GSFC/610.3)

ESMF-Enabled Coupled Models

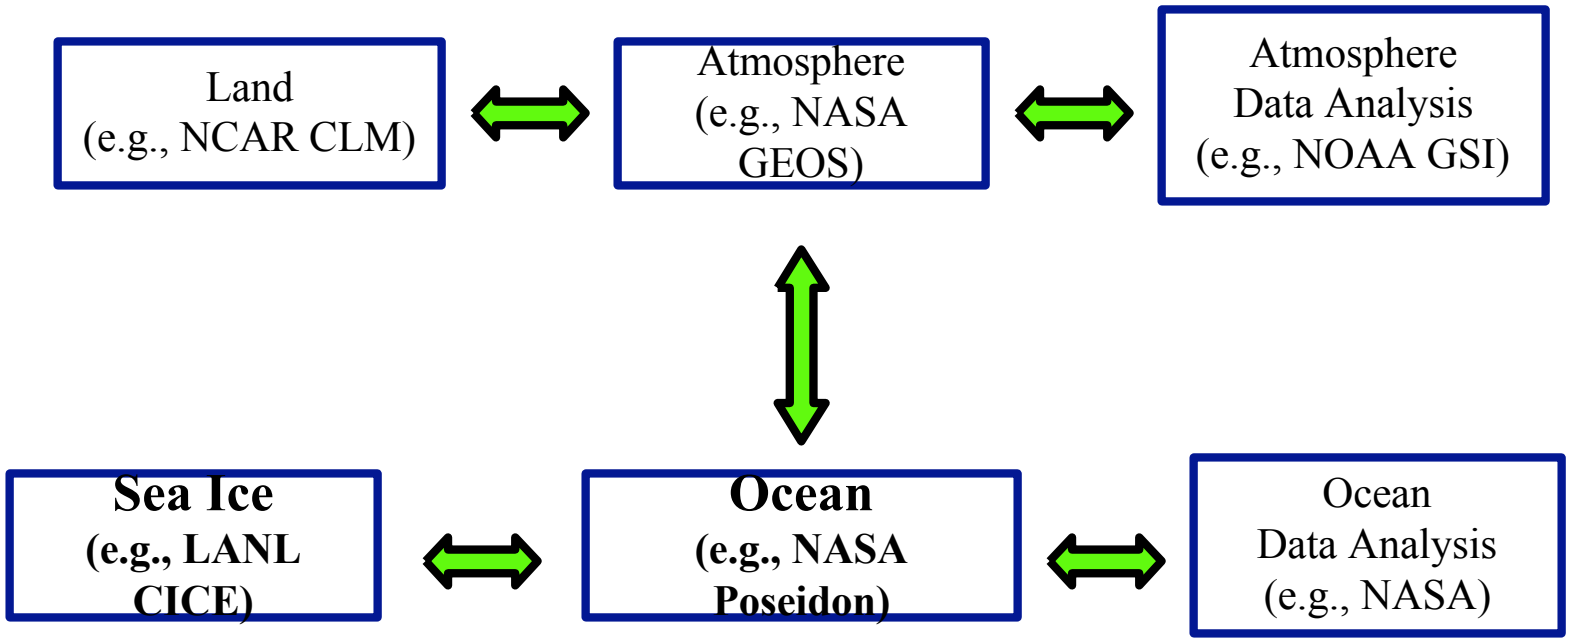

Example: Couple Different Ocean Models to NASA GEOS5

Source: Mike Seablom (GSFC/610.3)

DISTRIBUTED COMPUTING NODES

NASA Ames
Mt. View, California

Northrop Grumman
McLean, Virginia

NASA Goddard
Greenbelt, Maryland

NEXT-GENERATION NETWORKS

National Lambda Rail (10 - 40 Gb/s)

GEOS5/OGCM/ POSEIDON

GEOS5/OGCM/ MOM4

NCCS/NAS MODEL OUTPUT PORTALS

DAAC DATA PORTALS

SCIENCE ANALYSIS

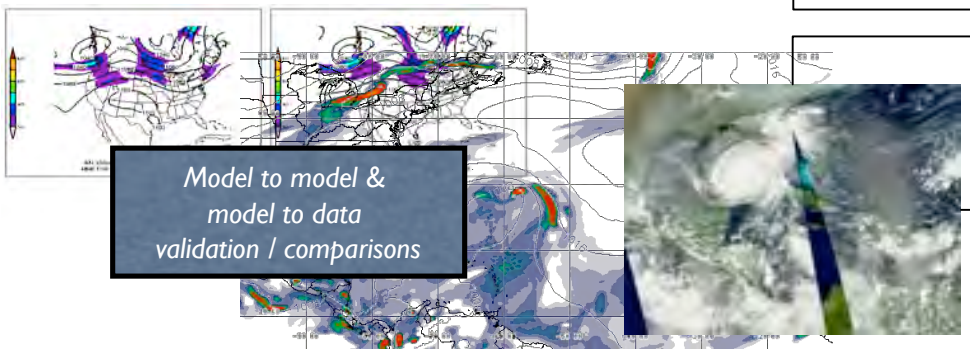

Model to model & model to data validation / comparisons

EXTERNAL COLLABORATORS

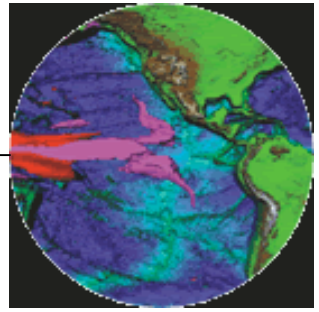

GFDL/NCDC MODEL OUTPUT PORTALS

GSFC's Various Uses of DRAGON
-- A Very Brief Overview --

Phoebus

Date: Tue, 1 Aug 2006 16:38:41 -0400
To: almes@internet2.edu
From: Pat Gary <Pat.Gary@nasa.gov>
Subject: Follow-up to Pat Gary(GSFC)'s interest in Phoebus

Guy,

In follow-up to my interest in Phoebus that I expressed to you during the 17Jul06 Jt Tech Wksp, please asap reply to the following.

1. Could we work with your Phoebus team to soon arrange another Phoebus test between your test host at ColumbiaU and a test host we'd provide at GSFC in Greenbelt, MD, where the Phoebus/HOPI ingress/eqress POP's would be NYC and DC (actually McLean, to which we connect via DRAGON)?

(Note: GSFC via its GISS (http://www.giss.nasa.gov/) has both an organizational and physical presence on ColumbiaU's campus. Hence a good Phoebus/HOPI test between NYC and DC could have potential "real world" value to us.)

…

As I mentioned to you during the 17Jul06 Jt Tech Wksp, my longer-term application for Phoebus is between GSFC and the Columbia supercomputer at ARC (as partly described in http://www.internet2.edu/presentations/jt2006jul/20060718-columbiascc-jones.pdf) via either HOPI LA-DC initially or NREN Sunnyvale-DC once you allow us to copy Phoebus….

"Brokering and Chaining Distributed Services and Data Using OptIPuter and the National Lambda Rail" by Ramapriyan (GSFC) et al to NASA's ROSES NRA

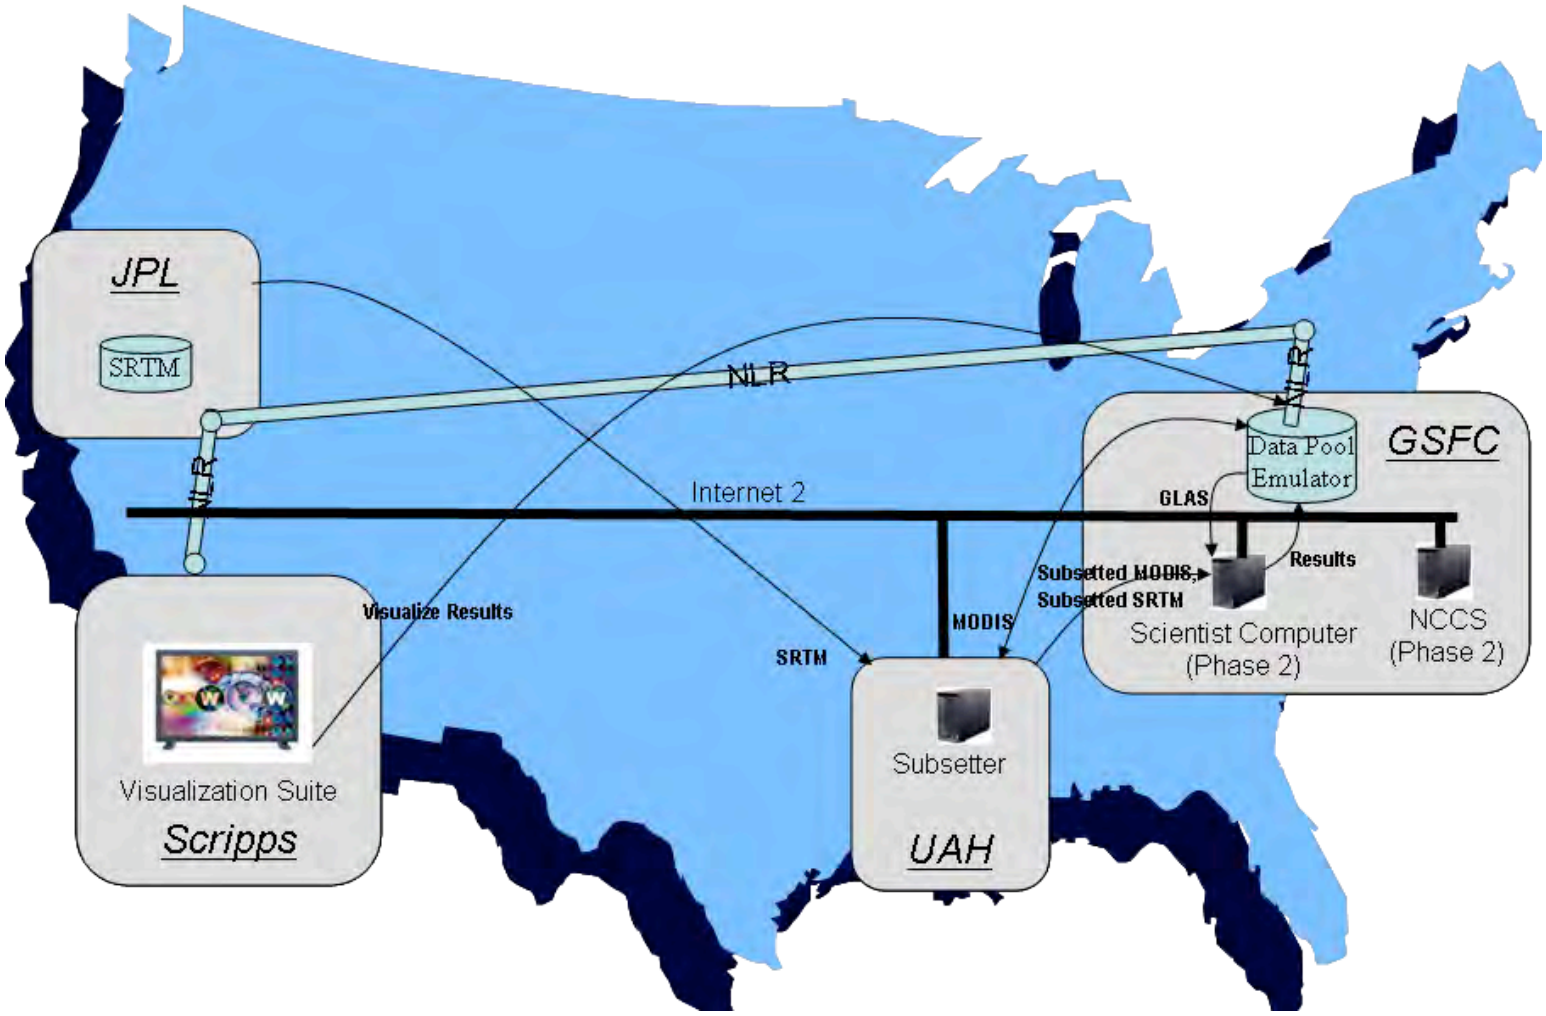

"Enabling NASA Applications Across Heterogeneous High Performance Networks" by Habib (CUNY) et al to NASA NNH05ZDA001N-Applied Information Systems Research (a.k.a. ROSES:D3)

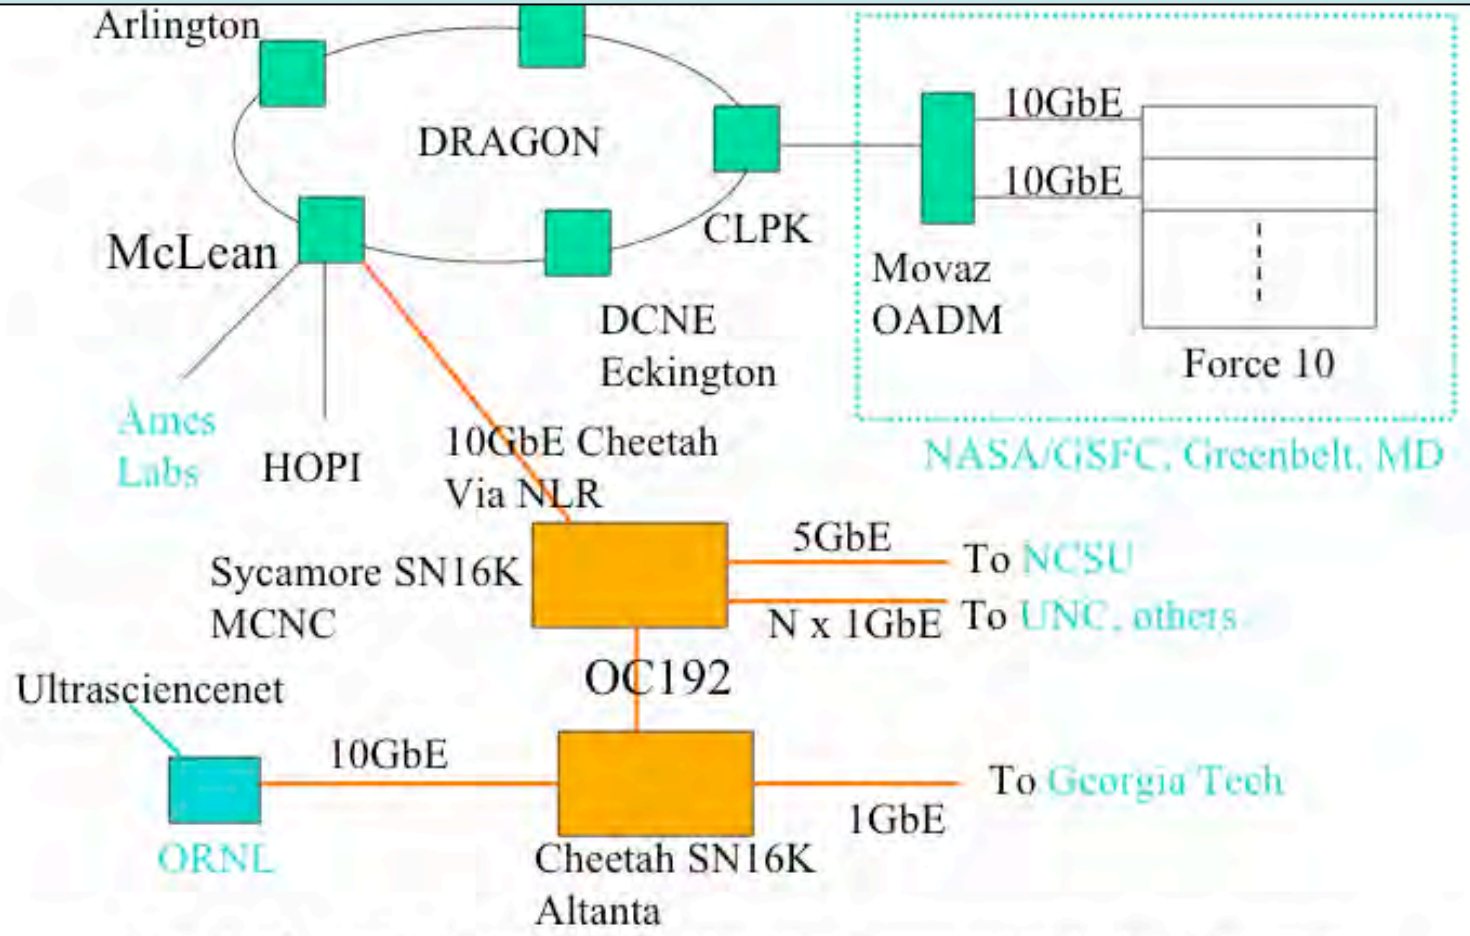

Fig. 1: Overall Proposed Network Connectivity to Cheetah

High Performance Networking and Remote Data Access
GSFC L-Net for NCCS and Science Buildings

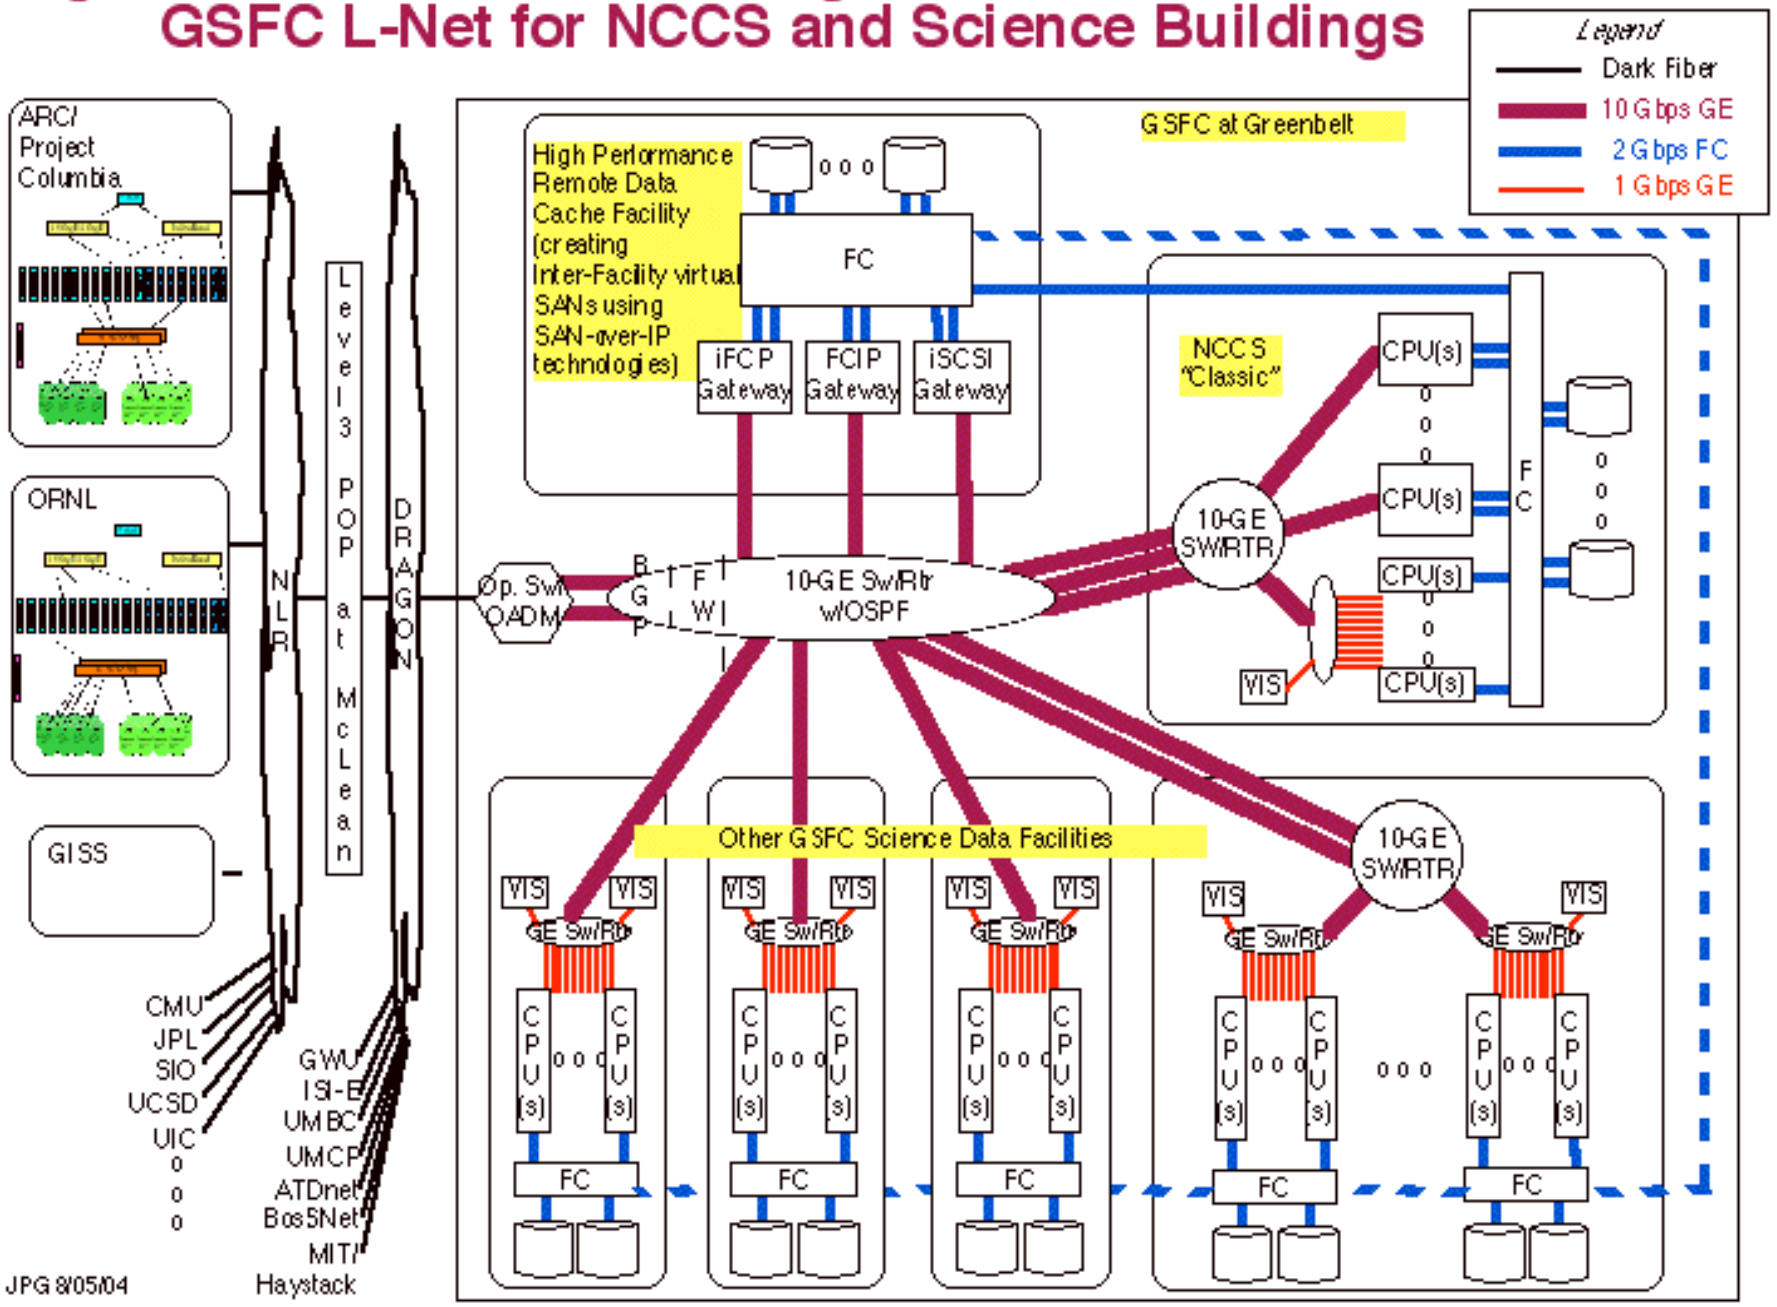

GSFC SEN and HECN

GSFC Lambda Network Project Website

- http://cisto.gsfc.nasa.gov/IRAD_Lambda.html
  - **Designs**
    - GSFC Local Network Part (i.e., within GSFC)
    - Regional Network Part (i.e., between GSFC in Greenbelt, MD, & Level3 POP in McLean, VA, typically involving the DRAGON optical network)
    - Transcontinental Network Part (i.e., use of NLR, GSFC 10-GE switch & workstations in the Level3 POP in McLean, VA, & remote end users/sites)
  - **Implementation Status**
    - GSFC Local Network Part
    - Regional Network Part
    - Transcontinental Network Part
  - **Presentations/Events in the News**
    - Eg: P. Gary's 18Feb05 presentation at GSFC's FY04 IRAD Colloquium <http://cisto.gsfc.nasa.gov/L-Netpdfs/FY04IRADGARY.pdf>
    - Live Demonstration of 21st Century National-Scale Team Science <http://www.calit2.net/articles/article.php?id=660>
  - **Related Links (e.g., DRAGON, HOPI, NLR, OptIPuter, …)**

GSFC SEN and HECN

Special Acknowledgements

GSFC Internal

- **High End Computer Network Team**
  - **Bill Fink/606.1**
  - **Kevin Kranacs/585**
  - **Paul Lang/ADNET/606.1**
  - **Aruna Muppalla/ADNET/606.1**
  - **Jeff Martz/CSC/606.2**
  - **Mike Steffenelli/CSC/606.2**
  - **Kevin Fisher/586/UMBC coop**
- **ESDIS Network Prototyping Lab**
  - **George Uhl/SWALES/423**
- **ESTC Computing Technology Project**
  - **PM: Jim Fischer/606**
- **IT Pathfinder Working Group**
  - **Chair: Dr. Milton Halem/Emeritus & UMBC**
- **Thunderhead Cluster**
  - **John Dorband/696**

GSFC External

- **National LambdaRail**
  - **CEO: Tom West**
  - **Net Eng Lead: Debbie Montano**
- **OptIPuter Project (NSF-funded)**
  - **PI: Dr. Larry Smarr/UCSD**
  - **Co-PI: Dr. Tom DeFanti/UIC**
  - **PM: Maxine Brown/UIC**
  - **UCSD Net Eng: Greg Hidley, Arron Chin, Phil Papodopolos**
  - **UIC Net Eng: Alan Verlo, Linda Winkler**
- **DRAGON Project (NSF-funded)**
  - **PI: Jerry Sobieski/UMCP**
  - **Co-I: Tom Lehman/USC-ISI/E**
  - **Net Eng: Chris Tracy/UMCP**
- **NASA Research and Education Network**
  - **DPM: Kevin Jones/ARC**